AD-A174 250 AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND ITS 1/2
POTENTIAL USE IN SPACE SYSTEMS(U) NAVAL POSTGRADUATE
SCHOOL MONTEREY CA G W MCDONALD JUN 86

UNCLASSIFIED F/G 22/2 NL

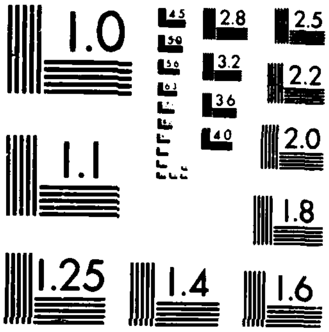

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

2

NAVAL POSTGRADUATE SCHOOL

Monterey, California

THESIS

AN INTRODUCTION TO ARTIFICIAL
INTELLIGENCE AND ITS POTENTIAL
USE IN SPACE SYSTEMS

by

Gary Wayne McDonald
June 1986

Thesis Advisor: Herschel H. Loomis, Jr

Approved for public release; distribution unlimited.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

AD-A174250

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION
UNCLASSIFIED

1b. RESTRICTIVE MARKINGS

2a SECURITY CLASSIFICATION AUTHORITY

2b DECLASSIFICATION / DOWNGRADING SCHEDULE

3 DISTRIBUTION / AVAILABILITY OF REPORT
Approved for public release; distribution is unlimited.

4 PERFORMING ORGANIZATION REPORT NUMBER(S)

5. MONITORING ORGANIZATION REPORT NUMBER(S)

6a. NAME OF PERFORMING ORGANIZATION
Naval Postgraduate School

6b. OFFICE SYMBOL *(If applicable)*
Code 32

7a. NAME OF MONITORING ORGANIZATION
Naval Postgraduate School

6c. ADDRESS *(City, State, and ZIP Code)*
Monterey, California 93943-5100

7b. ADDRESS *(City, State, and ZIP Code)*
Monterey, California 93943-5100

8a. NAME OF FUNDING / SPONSORING ORGANIZATION

8b. OFFICE SYMBOL *(If applicable)*

9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER

8c. ADDRESS *(City, State, and ZIP Code)*

10 SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO.	PROJECT NO	TASK NO	WORK UNIT ACCESSION NO

11 TITLE *(Include Security Classification)*
AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND ITS POTENTIAL USE IN SPACE SYSTEMS

12 PERSONAL AUTHOR(S)
McDonald, Gary W.

13a TYPE OF REPORT
Master's Thesis

13b TIME COVERED
FROM ________ TO ________

14 DATE OF REPORT *(Year, Month, Day)*
1986 June

15 PAGE COUNT
129

16 SUPPLEMENTARY NOTATION

17 COSATI CODES

FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS *(Continue on reverse if necessary and identify by block number)*

19 ABSTRACT *(Continue on reverse if necessary and identify by block number)*

This thesis provides an introduction to Artificial Intelligence and Space Systems, with comments regarding their integration. The survey of Artificial Intelligence (AI) is based upon a review of its history, its philosophical development, and subcategories of its current technologies. These subcategories are Expert Systems (ES), Natural Language Processing (NLP), Computer Vision and Pattern Recognition, and Robotics and Autonomous Vehicles. Emphasis is then directed toward the description of the fundamental characteristics of a generic space system system, including the space bus components, mission system components, ground node functions, and system missions. It is concluded that AI, in spite of its immaturity as a science, will prove to be a beneficial component of future space systems.

20 DISTRIBUTION / AVAILABILITY OF ABSTRACT
☒ UNCLASSIFIED/UNLIMITED □ SAME AS RPT □ DTIC USERS

21 ABSTRACT SECURITY CLASSIFICATION
UNCLASSIFIED

22a NAME OF RESPONSIBLE INDIVIDUAL
Herschel H. Loomis, Jr.

22b TELEPHONE *(Include Area Code)*
(408) 646-3214

22c OFFICE SYMBOL
Code 62Lm

83 APR edition may be used until exhausted
All other editions are obsolete

SECURITY CLASSIFICATION OF THIS PAGE

An Introduction to Artificial Intelligence and Its Potential Use In Space Systems

by

Gary Wayne McDonald
Lieutenant, United States Navy
B.A., Western Washington University, 1975
M.Ed. Western Washington University, 1976

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN SYSTEMS TECHNOLOGY
(SPACE SYSTEMS OPERATIONS)

from the

NAVAL POSTGRADUATE SCHOOL
June 1986

Author: Gary W McDonald
Gary Wayne McDonald

Approved by:
Herschel H. Loomis, Jr., Thesis Advisor

Elliot H Weinberg
Elliot H. Weinberg, Second Reader

Allen E Fuhs
Allen E. Fuhs, Chairman
Space Systems Academic Group

JN Dyer
John N. Dyer, Dean
of Science and Engineering

ABSTRACT

This thesis provides an introduction to Artificial Intelligence and Space Systems, with comments regarding their integration. The survey of Artificial Intelligence (AI) is based upon a review of its history, its philosophical development, and subcategories of its current technologies. These subcategories are Expert Systems, ~~(ES),~~ Natural Language Processing, ~~(NLP),~~ Computer Vision and Pattern Recognition, and Robotics and Autonomous Vehicles. Emphasis is then directed toward the description of the fundamental characteristics of a generic space system, including the space bus components, mission system components, ground node functions, and system missions. It is concluded that AI, in spite of its immaturity as a science, will prove to be a beneficial component of future space systems.

TABLE OF CONTENTS

I. AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE

Although artificial intelligence (AI) is one of the more glamorous "buzz words" in our current inventory, the word itself has been around for twenty-eight years, since it was first coined by John McCarthy [Ref. 1:p. 130]. The history of AI goes back at least 150 years to Charles Babbage, and if one wished to look, even farther.

In fact, if one applies a generous dose of imagination to the recitation of history the notion of artificial intelligence becomes an antediluvian one. Pamela McCorduck in her book, MACHINES WHO THINK, presents an interesting journey through history from which the following episodes are taken.

The Greek gods may be the first who were accomplished in the field of AI. As chronicled in the Iliad, it was Hephaestus, the god of fire and a divine smith, who built assistants for himself after being crippled:

> These are golden, and in appearance like living young women. There is intelligence in their hearts, and there is speech in them and strength, and from the immortal gods they learned how to do things [Ref. 2:p. 4].

However, before Hephaestus had set to his handy work, in about 1200 B.C., the inhabitants of the Siani were receiving an unmistakable message about dabbling in Godly affairs; "Thou shalt not make unto thee any graven image or any likeness of anything that is in heaven above or that is in the earth beneath, or that is in the water under the earth; Thou shalt not bow down thyself to them nor serve them..." Witness the effects of the transgressions of this message when God discovers the golden calf (Exodus: 32). But, the seed was planted, and infatuation with images, idols and robots was destined to be the blessing and bane of humans throughout history.

By 1580, it was not enough that a device should be made to think, but that it should embody human form. A rabbi, named Juden ben Loew, a contemporary and acquaintance of Tycho Brache and Johannes Kepler, is credited with the creation of an artificial man which he named Joseph Golem. Golem was fashioned from clay and brought to life by prayers and incantations, and by having the Holy Name impressed upon his forehead. Joseph was to serve Rabbi Loew in the capacity of a spy amongst the Gentiles who would occasionally rail against the Jews of Prague and bring harm upon them. When not forewarning Rabbi Loew of such uprisings, Joseph acted as a domestic servant in his house. As an example of the lack of specification in tasking this artificial man, Joseph was ordered to fetch water from the well and bring it into the house. The rub was that the amount of water was not specified and the entire content of the well was delivered. Over three centuries later a similar dilemma presented itself to Mickey when he tampered with the strength and magic of the Sorcerer's hat and found himself swamped by the efforts of the anthropomorphic broom.

In literature of the 1800's one finds androids as in Jean-Paul Richter's "The Death of an Angel," and machines taking on human traits as in "The Tales of Hoffman" by Offenback. Perhaps, the classic fictional example of the human desire to create a life like servant illustrated in the story of Frankenstein, from Mary Shelley's story of the same name (1818). Not only does Frankenstein stand as a story of human desire to artificially create one of its own, but more importantly and more poignantly Shelley artfully poses a series of moral and ethical questions about the use of science and the burden of responsibility which each scientist must carry for his creation. So impressed was Isaac Asimov with this burden of responsibility of the scientists, as they dabble in the artificial realm, that in 1950 he outlined the moral obligation of a robot in his "Three Rules for Robots":

1) A robot may not injure a human being or through inaction allow a human being to come to harm.

2) A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

3) A robot must protect its own existence as long as such protection does not conflict with the First and Second Law. [Ref. 3:p. 21]

Perhaps a whirlwind tour of history via a few fanciful stories provides little but entertainment; consider however the following points: first, it is from the extreme reaches of imagination coupled with scientific insight that come some of our most fruitful technologies, i.e., the flight of an airplane; second, the imagination demonstrated in these episodes allow us to examine many different concepts and consequences from a relatively safe environment; third, principles used in today's technologies may have been born in the very vivid imagination of yesterday's storyteller, wizard or tinkerman; and last, it is important that with any science its history can be formally and informally travelled to understand that new sciences don't simply open in full bloom one morning. Further, if a quick review of possible historical episodes seems garnished with a little trickery, slight of hand, incantation, and luck, this assumption is probably true--and true, not only in the past. If one were to query today's workers in the field of AI, the honest ones would affirm that this new science is filled with trickery, slight of hand, incantation, luck and even some efforts which represent progress!

Modern work in the area of AI probably began with Charles Babbage, and his unwavering supporter Countess Ada Lovelace. Babbage, a professor of Mathematics at Cambridge, planned for the invention of a machine called the Analytical Engine. His machine was to be mechanical, digital, use punchcards, and combine logic and arithematic processes to make logical decisions [Ref. 4:p. 22-28], [Ref. 5:p. 10], [Ref. 6:p. XVII]. Only portions of the Analytical Engine were ever built, but its plan did portend what was to take place 100 years later. The technology to construct Babbage's machine had simply not arrived.

In 1944, the first digital computer was built by H. Aiken, using about 300 electromagnetical switches [Ref. 6:p. XIX]. This marked the beginning of a technological era which would lead to the development of AI. A. R. Anderson, in his 1964 collection of papers, MINDS AND MACHINES, notes that "since 1950 more than 1000 papers have been published on the question as to whether "machines" can "think" [Ref. 7:p. 1].

This question is one which was, and still should be, examined in some philosophical detail by the AI researchers. Prior to the invention of the digital computer the question regarding the functioning of the human mind was given considerable discussion, but with the advent of the computer came a possible working model against which to compare theoretical ones. A statement of the topic may best be described by the following:

> We might say that human beings are merely very elaborate bits of clockwork, and that our having "minds" is

> simply a consequence of the fact that the clockwork is very elaborate, or we might say that any machine is merely a product of human ingenuity (in principle nothing more than a shovel), and that though we have minds, we can't impart that peculiar feature of ours to anything except our offspring: no machine can acquire this uniquely human characteristic." [Ref. 7:p. 2]

Or perhaps more simply, we are like computers, or they are like us, or we are unlike each other. If one of the first two cases is true, it may then be asked, "which came first, the revelation that our "mind" could be explained upon the model of a digital computer and we make a conscious effort to do so through their creation, or that the model of the computer is an unconscious organic outgrowth of the application of our mind?" These and similar questions kept busy the early (1950-1965) researchers in AI. Today, these philosophical questions remain unchanged by changing technologies. The approach to understanding them however has definitely changed.

The 1950's found researchers in AI extraodinarily excited and impressed by their newly developed device, the computer. Programs for the intelligent play of chess and checkers were produced, algorithms for carrying missionaries and cannibals across rivers in the same boat were lauded and the greatest effort of that time, the program 'General Problem Solver', (GPS) was written by J. C. Shaw, H. A. Simon, and A. Newell. [Ref.8]

GPS was an attempt to write a program which would "model human performance in search problems, such as puzzles and symbolic integration. Of course, not all problems can be thought of as search problems, so the 'G' in GPS seems a little optimistic now" [Ref. 9:p. 301]. Herbert Simon continues to work in the area of general problem solving with his latest programming effort, "Bacon."

Optimistic may be the best way to characterize the work of the 1950's to the mid 1960's. Optimism was so great that it prompted Marvin Minsky and Herbert Simon to claim in the mid sixties: "that within 20 years computers will be able to do everything humans can" [Ref. 10:p. 50]. In part, this optimism was fueled by the miraculous computational power that computers possess. However, if sheer strength in computation is to be the answer, the pioneers of the fifties would have been simply overwhelmed by the predictions as to contemporary computational power:

> Had a scientific prophet arisen in 1959 with the foresight to tell us we would, within 30 years, reach an age in which million-transistor ICs would enable the design of compact, parallel-processing computers capable of operating at rates of one trillion instructions per second (one million MIPS), he would have been ridiculed. Yet, today, we clearly see that this will be reality by 1989 [Ref. 11:p. 37].

And yet, it is not by the brute strength of computational power that advances in AI are made. It appears simply impossible for the computer to manipulate enough symbols which represent knowledge about the world to perform tasks as an intelligent person would do. The shortcoming is that "programs were missing crucial aspects of problem solving, such as the ability to separate relevant from irrelevant operations. . .They (MIT researchers) recognized that to solve "real world" problems the computer had to somehow gain real-world understanding and intuition." [Ref. 10:pp. 47-49]

A rather grim reality begin to dawn upon AI researchers as the 1960's came to a close: Computers and programs were not going to make impressive strides toward capturing any of the essence of human intelligence. It was Herbert Dreyfus in his 1970 book, WHAT COMPUTERS CAN'T DO, who delivered the coup de grace to this first era and perhaps inadvertantly breathed new life into the young science of AI. By 1970, the enthusiasm which characterized the 1960's had waned. The computer which would have the capabilities of the human being was not forthcoming. There was no general paradigm upon which to build "the thinking machine" which could replicate the human mind. So, the search began for approaches to AI based upon primary research from the earlier years, and for specification of problems which appeared to have a solution. Approaches to developing contemporary AI will be discussed later. It is important now to realize valuable lessons learned from 1950-1970. Two schools of thought seemed to prevail during this period, one which held that given enough facts and computational power a computer could perform as intelligently as a human. This simply hasn't yet been tested The second school held that all human intelligence processes can be subdivided such that if these sub-divisions are small enough they could be modeled using the computer. Unfortunately, this tends to create subdivisions which resemble "toy problems," i.e. trivial problems with overly simple solutions. Although the solutions to these toy problems may be simple, their solutions are unlikely to represent the real world either singly or collectively. Thus, a new avenue of exploration had to be opened.

Realizing that the replication of human intelligence, in total, was for the time an impossibility, AI researchers of the 1970's set about defining specific tasks or problems areas generally addressed by human intelligence and brought the power of computers to bear upon them. These areas, which are generally accepted as those making up the specific field of AI are; computer vision and image understanding, expert systems, natural language processing, and robotics and autonomous vehicles.

From these categories research developed which was domain specific, objectifiable, whose solution had other than trivial purposes, and which could be demonstrated. The 1970's offered expert systems such as XCON, DIPMETER ADVISOR, INTERNIST and SOPHIE. Advances in computer vision and image understanding which enable a cruise missile to find its way to its target via a TERCON guidance system (terrain and contours guidance system) were developed. Natural language processing was created which allows language understanding. Robotic devices, while limited in their ability when compared to humans, are working on production lines world-wide.

The 1970's was a time of optimizing the type of problems to which AI techniques could be applied and of developing and demonstrating these solutions. In the 1980's this method of objectifying the problem domain continues and progress is slow.

Before proceeding to greater specificity in this discussion, it is important to develop a definition of what AI is, and to do so means, in many regards, to approach it philosophically. This takes one back to the early experimenters in AI beginning with A.M. Turing who argues the following: "The reader must accept it as a fact that digital computers can be constructed, and indeed have been constructed according to the principles we have discussed, and that they can, in fact, mimic the actions of a human computer very closely" [Ref. 5:p. 9]. Turing fails to define the scope in which the digital computer might mimic the human computer, or how well this mimicing might be demonstrated. However, one may still turn to his writings, and those of others, in order to refine this question.

Feigenbaum and Feldman from their book, COMPUTERS AND THOUGHT, put the question squarely before us: "Is it possible for computing machines to think?"

> No-if one defines thinking as an activity peculiarly and exclusively human. Any such behavior in machines, therefore, would have to be called thinking-like behavior.
>
> No-if one postulates that there is something in the essence of thinking which is inscrutable, mysterious, mystical.
>
> Yes-If one admits that the question is to be answered by experiment and observation, comparing the behavior or the computer with that behavior of human beings to which the term "thinking" is generally applied [Ref 12:p. 3].

The affirmative alternative is selected and the authors state their answer to the question as the goal of AI research: "to construct computer programs which exhibit behavior that we call "intelligent behavior," were we to observe it in humans. [Ref. 12:p. 3]

One might suggest Feigebaum and Feldman would agree that "they (digital computers) can in fact mimic the actions of a human computer very closely" [Ref. 5:p. 9]. So, it seems reasonable to conclude that Turing, Feigebaum and Feldman would define computing machines to be "thinking machines."

H. I. Dreyfus in his book, WHAT COMPUTERS CAN'T DO, points to another conclusion about thinking machines. First, sighting from a paper by Newell and Simon:

> It can be seen that this approach (information processing) makes no assumption that the "hardware" of computers and brains are similar, beyond the assumptions that both are general purpose symbol manipulating devices, and that the computer can be programmed to execute elementary information processes functionally quite like those executed by the brain [Ref. 6:p. 67].

He then calls into question the definition of a general purpose symbol-manipulating device, and elementary information processes. Dreyfus addresses these questions:

> The assumption that humans function like a general-purpose symbol manipulating device amounts to:
>
> 1) a biological assumption. . .that the brain processes information in discrete operations by way of some biological equivalent of on/off switches.
>
> 2) a psychological assumption that the mind can be viewed as a device operating on bits of information according to formal rules. . .the computer serves as a model of themind as conceived by empiricists such as Hume or idealists such as Kant. Both. . . have prepared the ground for the model of thinking as data processing - a third person process in which the involvement of the "processor" plays no role.
>
> 3) an epistemological assumption that all knowledge can be formalized in terms of Boolean functions and the logical calculus which governs the way the bits were related according to rules.
>
> 4) an ontological assumption. . .since all information fed into digital computers must be in bits, the computer model of the mind presupposes that all relevant information about the world, everything essential to the production of intelligent behavior, must in principle be analyzable as a set of

> situation-free determined elements. . .that there is a set of facts each logically independent of all the others [Ref. 6:p. 68].

The biological assumption he concludes "is an empirical hypothesis which has had its day" [Ref. 6:p. 74]. There exists no empirical evidence that we can duplicate in a machine the bio-chemical nature of the brain.

As for the psychological assumption even though the school of behavioralistic psychology seeks to explain human intelligence processes in terms of man's behavior, there is again a lack of empirical evidence that even if his behavior could be described as sets of rules governing his actions, that his behavior could be duplicated by a machine.

Dreyfus succinctly states the refutation of the epistemological assumption.

> A full refutation to the epistemological assumption would require an argument that the world cannot be analyzed in terms of determinate data. Then, since the assumption that there are basic unambiguous elements is the only way to save the epistemological assumption from regress of roles, the formalist, caught between the impossibility of always having rules for the application of the rules and impossibility of finding ultimate unambiguous data would have to abandon the epistemological assumption altogether. [Ref. 6:p. 117]

The contradiction to the ontological assumption is in the fact that learning intelligent behavior lies in generality and flexibility and that in the end, the whole problem may never have been hidden by a set of facts independent of all others, but that a gradual recognition of the whole is the synergism of the many separate sets. In other words, relevant facts can not be separated from the significance of the whole while the whole remain unblemished and clearly defined.

Dreyfus takes his stand in what might be called the realist's approach to the problem which begins to define AI in terms of unique application of computers to well specified problems, a theme which pervades the 1970's and 1980's. Dreyfus concludes, "It no longer seems obvious that one can introduce search heuristics which enable the speed and accuracy of computers to bludgeon through in those areas where human beings use more elegant techniques...only newer and faster machines, better programming languages, and cleverer heuristics can continue to push back the frontiers" [Ref. 6:pp. 138-139]. Hardly an enthusiastic endorsement for the thinking machine.

The argument of the "thinking machine" remains in the gray region, as Scriven suggests:

> There appears to be a paradox associated with the concept of a conscious machine. On the one hand, it does not seem that there is anything in the construction, constituents, or behavior of the human being which it is essentially impossible for science to duplicate and synthesize. On the other hand, there seems to be some important and meaningful descriptions of human behavior which can never be properly applied to machines. We feel puzzled that the basis for a description can be reproduced, yet the description cannot be reapplied. [Ref. 13:p. 35]

This holds no promise to clearing up the confusion surrounding the "thinking machine" so the debate continues.

Turing develops his argument for the "thinking machine" by refuting the following objections:

1) The Theological

2) The "Heads in the Sand"

3) The Argument from Consciousness

4) The Arguments from Various Disabilities

5) The Mathematical Objection

6) Lady Lovelace's

7) The Argument from Continuity in the Nervous System

8) Argument from Informality of Behavior and

9) The Argument from Extrasensory Perception

[Ref. 5:pp. 14-24].

With all of this it has not been decided if machines can "think" nor a definition of AI supplied.

In their book, ARTIFICIAL INTELLIGENCE THROUGH SIMULATED EVOLUTION, Fogel, Owens, Walsh offer a definition of AI by first drawing upon the words of Lord Kelvin:

> I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is a meagre, an unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science, whatever the matter may be [Ref. 14:p. 1]

Now to their definition. . ."knowledge being the useful information stored within the individual and intelligence being the ability of the individual to utilize the stored information in some worthwhile (goal directed) manner" [Ref. 14:p. 1].

Then from this definition, these authors led to the "Avenues Toward Artificial Intelligence" of which there are three from which to choose. The first, for which the term bionicism is applied, is that which views the human as the highest life form and worthy of emulations by machine. Bionigists, as others, also hold little hope that the brain may be mechanically modeled.

The second group, heuristicists are those who would define AI as a set of rules which govern behavior and produce an appropriate response.

And thirdly, "from a less egocentric standpoint, the human animal may be viewed as but a simple artifact of the natural experiment called evolution. Though, certainly, man is an intelligent creature, there is no reason to believe that he is the most intelligent creature that could possibly exist...In a sense, the evolutionary approach asks, "what might intellect be like sometimes in the distant future?...the evolutionary approach...is an attempt to model processes as they might occur in nature; to describe what ought to be rather than what is" [Ref. 14:pp. 3-10). No small task and they do offer an interesting if not grandiose definition of AI.

> Artificial intelligence is realized only if an inanimate machine can solve problems that have, thus far, resisted solution by man, not because it can discover for itself new techniques for solving the problem at hand [Ref. 14:p. 8].

Most would be satisfied if a machine could simply assist them in their decision making processes, a task for which machines are well suited.

Others (McCune and Drazovich, Winston, and Dreyfus) offer their definitions of AI but perhaps a working definition is supplied by Manuel and Evanczuk:

> AI. . .is the part of computer science concerned with designing intelligent computer systems; that is, systems that exhibit the characteristics associated with intelligence in human behavior, such as understanding, learning, reasoning and problem solving. [Ref. 1:p. 128]

This broad definition may be narrowed to say the "Artificial Intelligence is a field defined by its objectives..." [Ref. 15:p. 273]. And, this may very well be the way it must be programitically defined for the time being. Our only other alternative is to conclude as Kolbus does "Even experts in the field cannot decide on an exact definition of artificial intelligence" [Ref. 16:p. 97].

As illusive as a definition of AI appears to be, an operational definition is offered for the sake of common ground for further discussion. And that is, AI is the science of programming computers in a manner in which their output, and perhaps in some ways their internal processes, reflect, resemble, and take on the characteristics of human thought processes, and the decisions reached by these processes. The mood of contemporary researchers in the field of AI may be away from the use of the confusing term artificial intelligence to describe the science and to a more restrictive, and more easily defined term such as computer aided decision support systems. For, it is not until the computer can automatically evaluate and update its data base, and create and delete rules, and change the priority and firing order of its rules that learning occurs, and AI is achieved in a genuine sense.

Short of a machine acquiring the ability to learn, how does one program a computer to act in what appears to be an intelligent manner?

Intelligent behavior may often be associated with problem solving and it is here that AI programming can achieve some success. Problem solving strategies may differ in their details but for the most part they consist of the following steps:

1) description of the current situation or problem (often referred to as an initial state) to be solved and a description of the solution (goal state)

2) evaluation of the problem, i.e. is it simple or complex? If it is complex can it be decomposed into a simpler problem for solution?

3) evaluation of the resources which may be brought to bear upon the problem (often referred to as operators, i.e., rules, heuristics, symbolic manipulators).

4) assignment and application of the appropriate resources which would offer a solution to the problem.

5) evaluation of the outcome of step four to assertain the correctiveness of the solution or the partial or full attainment of the goal state.

6) iteration or recursion of the steps until a "better solution" appears impossible to achieve.

Clearly, in step one, the problem must be of the family of problems for which the program is designed to solve. Once this criterion is met there are basically two methods of approach with which to proceed, a brute force method and the tailored method of solution. The brute force method employs a technique which evaluates every possible state and a combination thereof, based upon its data base, to achieve a solution. This approach is most advantageous when applied to a simple problem where the solution is heavily dependent upon straight forward mathematical calculation, and an iterative or recursive program may be applied. The brute force method reveals its weakness when the problem space becomes more complex. For example, a brute force program is inappropriate for playing chess:

> In chess, the average number of moves that can be made from a given position is 35; an exhaustive search only three moves deep for each player would require the examination of more then 1.8 billion moves. [Ref. 18:p. 119].

Finally, is a brute force program not much more than a straight forward algorithm, and should such a program be described as AI? This type of programming hangs on the fringes of AI, and may often be marketed as AI.

It should also be noted that a brute force approach may not simply rely upon an exhaustive search of data space utilizing a single criterion for achieving a goal.state, but may include a series of rules which fire sequentially and repeatedly until this series has exhaustively evaluated all states to achieve the goal state. This use of rules in a sequential manner should not be confused with heuristics or rule of thumb presented in the tailored approach. The use of the word rule here is used as an operator whereas rule (heuristic) may be both an operator and evaluator. The tailored search may be most easily described as follows:

> in most artificial intelligence programs heuristic principles or informal rules of thumb, are incorporated so that the most promising actions are selected early into the search and less promising ones are eliminated from full scale consideration" [Ref. 17:p. 118]. For example, ". . .virtually all chess programs devised since the 1950's are based on a heuristic search. This rule of thumb technique is incorporated because a brute force search alone simply does not pay off" [Ref. 17:p. 119].

Turing gives his explanation of heuristic search as a "rule of thumb, strategy, trick, simplification, or any other kind of device which drastically limits search for solutions in large problem space" [Ref. 5:p. 6]. Thus, the heuristic approach allows the computer to narrow the field of view quickly without consideration of every possible alternative but concentrating on those relevant to the question at hand.

Thus, the heuristic (tailored) method may be seen to be a more sophisticated approach to a solution and perhaps the only way to proceed through a complicated problem space. However, it too has its shortcomings, the most predominant being the level of effort required to write and maintain the program. This level of programming effort compared to its output may prove costly in every regard and make it an unattractive choice for problem solving. (Suitable problem spaces will be discussed later under the topic "expert systems"). However, as Elaine Rich points out in her book. ARTIFICIAL INTELLIGENCE:

> Heuristics are like tour guides. They are good to the extent they point in interesting directions; they are bad to the extent that they lead into dead ends. Some heuristics help to guide a search process which might previously have been overlooked. Others (in fact, many of the best ones) may occasionally cause an excellent path to be overlooked. But, on the average, they improve the quality of the paths that are explored. Using good heuristics, we can hope to get good (even if nonoptimal) solutions to hard problems, such as the traveling salesman. [Ref. 18:p. 35]

Rich later cites H. Simon:

> Rarely do we actually need the optimum solution; a good approximation will usually serve very well. In fact, there is some evidence that people, when they solve problems, are not optimizers but rather are satisfiers [Ref. 18:p. 36].

Why should one be satisfied with a less optimal solution? The strongest argument in its favor may be that as problem spaces become complicated a brute force search becomes a combinatorial explosion (exponential in growth).

Now that two general strategies to problem solving have been described, more specific features of searching may be briefly examined. The specific features or tools are breadth first search, depth first search or backtracking, (chronological backtracking) and means-ends analysis. Utilizing a decision tree, a breadth-first search is undertaken as an exhaustive exploration of each state at each level of the tree until a suitable goal state is reached.

> Breadth-first search corresponds to always putting new states on the end of the queue, that is, managing it "first in/first out". (Thus) it examines all states that are in operator applications from the initial state before any that are N+1 away. [Ref. 9:p. 266].

On the other hand, depth first search will embark on exploration on a specific branch of the decision tree and follow that branch until the branch is exhausted, at which time the search is again initiated upon another branch, or until a suitable goal state is achieved, or a depth search or level limitation is envoked.

> Depth-first search, or backtracking, corresponds to putting new states on the front of the queue, that is, managing it "last in/first out". This means that if two states S1 and S2 are produced by applying operators to a state S, then every state reachable from S1 will be examined before any reachable from S2 (unless some are reachable from both). [Ref. 9:p. 266].

evaluation of previous states searched in a branch without having to conduct the search from the initial state.

And lastly, the means-ends analysis which was developed and programmed by Newall, Shaw, and Simon in their General Problem Solver (GPS), which, as its title suggests attempts a genuinely global approach to problem solving rather than a domain specific problem.

> The means-ends analysis first determines the difference between the initial and goal states and selects the particular operator that would most reduce the difference. If this operator is applicable in the initial state, it is applied and a new intermediate state is created. The difference between this new intermediate state and the goal state is then calculated and the best operator to reduce this difference is selected. The process proceeds until a sequence of operators is determined that transforms the initial state into the goal state.
>
> The difference reduction approach assumes that the differences between a current state and a desired state can be defined and the operators can be classified according to the kinds of differences they can reduce. If the initial and goal states differ by a small number of features, and operators are available for individually manipulating each feature, then difference reduction works. However, there is no inherent way in this approach to generate the ideas necessary to plan complex solutions to difficult problems. [Ref. 19:p. 26]

Throughout this entire discussion an underlying assumption has been made, which is that the computer has knowledge of a particular domain and that this knowledge may have an adequate representation so that it may be evaluated by the computer. The literature offers overwhelming support for knowledge as the foundation of AI.

knowledge of a particular domain and that this knowledge may have an adequate representation so that it may be evaluated by the computer. The literature offers overwhelming support for knowledge as the foundation of AI.

> To make a program intelligent, provide it with lots of high-quality specific knowledge about some problem area. [Ref. 19:p. 4]

As Dreyfus and Dreyfus cites from Marvin Minsky's SEMANTIC INFORMATION PROCESSING:

> . . .I therefore feel that a machine will quite critically need to acquire on the order of a hundred thousand elements of knowledge in order to behave with reasonable sensibility in ordinary situations. A million, if properly organized should be enough for a very great intelligence. [Ref. 10:p. 49]

But this attractive pool of knowledge which is required for an intelligent system is not with out problems of its own:

> One of the few hard results to come out of the first 20 years of AI research is that intelligence requires knowledge. To compensate for its overpowering asset, indispensability, knowledge also possesses some less desirable properties; including:
>
> - It is voluminous
> - It is hard to characterize
> - It is constantly changing [Ref. 20:p. 5]

Further testimony:

> But as the amount of knowledge grows, it becomes harder to access the appropriate things when needed, so more knowledge must be added to help. But now there is even more knowledge to manage, so more must be added, and so on. [Ref. 20:p. 21]

Thus, given that a program has enough knowledge to be intelligent, how will this knowledge be utilized? Obviously knowledge must be manipulated in order to draw conclusions about and offer solutions to a problem, but prior to this manipulation the knowledge must be presented. There are three basic activities that must occur:

1) knowledge of the world must be extracted from the world and stored as a representation or model of the world in the computer.

2) knowledge stored in the computer must be internally represented in such a way as to allow easy accessibility and operation on it.

3) internally stored knowledge must be translatable and presented in a manner useful to human beings .

How is this representation to take place? Even at its most simplistic, elements of the representation of knowledge become very complicated, very quickly.

> Representation schemes are classically classified into declaration and procedural ones. Declarative refers to representation of facts and assertions, while procedural refers to actions, or what to do. It is virtually impossible to come up with a pure system or either type as ultimately both assertions and what to do with or about them are involved in the data structures and the access mechanism in any knowledge representation.
>
> A further subdivision for declaration (objective oriented) schemes includes relational (semantic network) schemes and logical schemes. [Ref. 19:p. 201]

This has merely divided knowledge into gigantic domains which must undergo innumerable subdivisions in order to make knowledge useful to the program. These subdivisions of representation, manipulation are provided for by the predicate calculus:

> . . .the predicate calculus. . .gives us a way of calculating the truth of propositions. As such, the predicate calculus consists of a language for expressing propositions and rules, or how to infer new facts (propositions) from those we already have. [Ref. 9:p. 15].

But now the problem of usefully representing the world seems to grow exponentially again. How is the computer to "know" when is a chair a table and not a chair? These seemingly obvious facts are easily understood and manipulated by humans every day, and yet, it is with painstaking difficulty that obvious facts are made useful by the computer. The process is called common sense programming and it doesn't exist today. Or as Hubert and Stuart Dreyfus put it:

> To explain our own actions and rules, humans must eventually fall back on everyday practices and simply say "This is what one does". In the final analysis, all intelligent behavior must hark back to our sense of what we are. We can never explicitly formulate this in clear cut rules and facts, therefore, we cannot program computers to process that kind of know-how. [Ref. 10:p. 51]

Upon even a cursory inspection of the field of AI one is likely to conclude, and quickly, that creating AI is a difficult business. And, as would be expected if something is difficult to create, it is also expensive.

Given a problem domain there is a rather simple heuristic which would indicate trends in the cost of the production of an AI system.

> There are two opposing ways to improve the efficiency (solution time) of a problem solver:
>
> *use a cheap evaluation function and explore lots of paths that might not work out, but in the process acquire information about the interrelationships of the actions and the states as an aid in efficiently guiding a subsequent search.
>
> *use a relatively expensive evaluation function and try hard to avoid generating states not on the eventual solution path. [Ref. 19:p. 27]

What does this type of expertise cost? Davis reports that "developing a substantial expert system with real performance takes at least five years of effort, assuming the team already has some background in AI problem solving techniques. If the team is starting from scratch, with this technology, then developing a high performance expert system can take considerably longer" [Ref. 21:p. 26].

The next question is likely to be "what does it cost in dollars and cents?" Again, Davis answers:

> it is not difficult to find real problems where an expert performs slightly better than the average person doing the same job and where the disparity is extremely costly. . .at times the benefit of simply narrowing this gap can range into terms of millions of dollars per year. Clearly, the economic consequences of the technology are substantial [Ref. 21:p. 37]

Manuel and Evanczuk report an "estimated $66 million to $75 million" was spent for AI in 1983. They go on to state that "International Resource Development Inc. of Norwalk, Conn., predicts an estimated U.S. market for AI products and services of $66 million in 1983, growing to $8.5 billion by 1993" [Ref. 1:p. 127].

For the answers to the questions: Where do these people and this money come from? we need only turn to the experts once again. Davis reports that, "of approximately 2,500 people actively working on AI in the United States, fewer than 250 are experienced and actively working in the area of expert systems" [Ref. 21:p. 38]. Marvin Minsky is even more pessimistic, "The number of people doing basic research in AI is probably under one hundred people and may be under fifty". Further, Minsky states that, "there is no significant increase in the number of people working on ideas that we will want to use in ten years" [Ref. 22:p. 295]. "In the most optimistic counting, those universities produce two or three PhD's a year. Yale gets a representatively good sample of new graduate students, the kinds that had all A's from Ivy League schools; 50 percent of thos^e^ students fail, never getting PhD's because the work is too hard. It requires a mind they do not have, a certain level of imprecision, and an ability to retain a tremendous amount of knowledge. Students work on massive projects that take two to three years to complete. Adding resources does not help. If you put ten people on the same project, it does not go ten times faster. It may go ten times slower" [Ref. 23:p. 147]. To conclude, "acquisition of good people is one of our largest problems...simply because there are not enough good people to go around". [Ref. 24:p. 91]

Countries other then the United States are also committed to AI research. The Japanese are a few years into a 10 year fifth generation computer project. The U.K. has Alvery, with sponsorship by ICU, PLC, Sinclair Research, THRON-EMI, Shell Oil Co. Labs, and two government agencies. Schlumberger, located in Paris, makes the largest contribution to the French effort, and the European Commission has the Espirit project [Ref. 1:p. 129]

Facing these monumental difficulties people continue seeking applications in all fields of endeavor. "The defense Science Board named AI and robotics as one of the technology areas with the greatest potential for the DoD" [Ref. 25:p. 87] Meyrowitz reports "a strong argument can be made that the military applications of AI offer the greatest challenge to the science" [Ref. 26:p. 45].

The educational community strives for new developments in machine aided learning. The commercial community awaits the Intelligent Computer with the basic consideration that "Theoretically there is no task to which an expert computer could not be assigned" [Ref. 16:p. 98].

AI is here to stay. Its needs are well-trained people, and precise workable problems. Through its ups and downs, "No one this time expects AI to revert to its former status as an academic curiosity...for many very practical people, AI is no longer science fiction; like space, it has begun to be part of the real world" [Ref. 27:p. 9].

This brief introduction to AI is intended to serve as a stimulus to further exploration by the reader and to prove beneficial to understanding the remainder of this paper. The next four chapters will undertake the description of the four topical areas which are identified earlier as making up the body of AI. There are, again expert systems, a natural language processing, computer vision and scene recognition, and robotics and autonomous vehicles.

It cannot be overstressed that artificial intelligence is the creative programming of computers in a fashion that when they operate on a problem, they do so in such a way as human intelligence might operate on that same problem. And, that by the operation of these creative programs they may aid humans in understanding problem spaces and ultimately assist in problem solving. There are no black and which answers in this new science of AI and all who are so inclined, regardless of their background, are engouraged to indulge themselves in it.

II. EXPERT SYSTEMS

Expert systems (ES), as stated earlier in this paper, are a subcategory of AI whose origins may be traced to at least three areas: "symbolic programming, cognitive psychology and work on incremental programming environments" [Ref. 28:p. 52]. Although initial efforts to develop an ES are attributed to the Stanford group, who developed DENDRAL in the 1960's, the commercialization of ES began in 1980 and 1981 [Ref. 28:p. 52]. Today's ES may be defined as a sophisticated computer program, which utilizes representations of expert human knowledge, via logical symbolic manipulation, to solve problems generally solved by humans, within a very limited problem space.

Knowledge is the key to the emerging ES, rather than formal reasoning, for several reasons:

- most difficult and interesting problems do not have tractable algorithmic solutions. . .
- human experts achieve outstanding performance because they are knowledgeable.
- in short, an expert's knowledge per se seems both necessary and nearly sufficient for developing an expert system

[Ref. 29:pp. 3-5].

It is, then, knowledge which is the underpinning of the ES and distinguishes it from more simple algorithmic programming and data manipulation. Donald A. Waterman, in his book A GUIDE TO EXPERT SYSTEMS, offers the following comparison:

Data Processing	Knowledge Engineering
Representations and use of data	Representation and use of knowledge
Algorithmic	Heuristic
Repetitive Process	Inferential Process
Effective manipulation of large data bases	Effective manipulation of large knowledge bases

[Ref. 20:p. 24]

And finally, Brackman, et al., notes that "In sum, the first important factor that distinguishes work on expert systems from simply high-quality, special-purpose programming is its relation to AI in general and to symbolic representational reasoning in particular". [Ref.30:p. 46]

With knowledge and the ability to reason symbolically what then can one expect on ES do? ES has been developed in many areas including Chemistry, Medicine, Geology, Mathematics, Computer Systems, etc. Categorically speaking however, an ES is perhaps best suited for the following:

Interpretation
Design
Diagnosis
Control
Debugging
Repair
Prediction
Instruction

[Ref. 20:pp. 42-48]

Thus, within any category there exist problems which are potentially well suited for solution with an ES; however, any single problem may be so complex as to become very quickly overwhelming. The first criterion for the selection of an appropriate problem to attack with an ES is a narrow and specialized domain. Not only is it required that a problem space be highly specified but that data and knowledge relating to the problem should be reliable and static. [ref. 31:p. 92]

Once a good problem has been defined, two questions must be answered: first, when should an ES be used and, secondly, why should an ES be used?

To answer the first, an ES should be used when:

- Task does not require common sense·
- Task requires only cognition skills
- Experts can articulate their methods
- Genuine experts exist
- Experts agree on solutions
- Task is not too difficult
- Task is not poorly understood

[Ref. 20:p. 128]

And, for the answer to the second question, why should an ES be used:

- Task solution has a high payoff

- Human expertise (is) being lost
- Human expertise (is) scarce
- Expertise (is) needed in many locations
- Expertise (is) needed in hostile environments

[Ref. 20:p. 130]

E. A. Feigenbaum considered other reasons for justifing the use of an ES in his paper, "Knowledge Engineering: The Applied Side". Can an ES reduce cost and save time? Perhaps with development and maintenance costs of an ES being high today, its use in many situations may not be justified. However, Feigenbaum argues that inevitably, the cost of an ES will fall as the cost of computers has fallen and "computers will act as intelligent assistants" to professionals. The most important gain seen by this Feigenbaum is "The gain to human knowledge by making explicit the heuristic rules of a discipline will perhaps be the most important contribution of the knowledge-base systems approach" [Ref. 32:pp. 49-50].

So, once a problem has been chosen and the justification for the development of an ES has been made, the difficult work begins; building the ES. This task can be divioed into the following subtasks:

- Finding the required knowledge (knowledge acquisition)
- Representing the knowledge in the computer (knowledge engineering)
- Constructing the inference engine

- Understanding conclusions
- Explaining the conclusions

There can be little debate that the "The accumulation and codification of knowledge is one of the most important aspects of an expert system" [Ref. 20:p. 7], but where does the knowledge come from? Knowledge, in an area of expertise, can be essentially divided into two categories, that which is factual and that which is heuristic in nature. Factual information is widely available in text books and journals. Heuristic knowledge is more difficult to ferret out:

> This is the knowledge which constitutes the rules of expertise, the rules of good practice, the judgemental rules of the field, the rules of plausible reasoning. These rules collectively constitute what the mathematician, George Polya, has called the "act of good guessing." In contrast to the facts of a field, its rules of good guessing are rarely written down. This knowledge is transmitted in internships, PhD. programs, apprenticeships. [Ref. 32:pg. 37]

Obviously, factual knowledge is rather straightforward and easy to extract from the problem domain but, where does one mine hueristic knowledge? Weiss and Kulikowski believe many types of information may be supplied by the expert describing:

- personal experience of past problems solved
- personal expertise or methods for solving the problems
- personal knowledge about the reasons for choosing the methods used.

[Ref. 33:p. 11]

Waterman adds to the list:

- on site observation
- problem discussion
- problem description
- problem analysis
- system refinement
- system examination
- system validation

[Ref.20:p. 188]

It must be understood that these lists are merely suggested as starting points for beginning the time-consuming extraction of heuristic knowledge from experts. This process of knowledge acquisition is not only time consuming, but perhaps the most critical aspect of any ES, for never has the adage "garbage in - garbage out" been more appropriate. Correct, concise, understandable, complete knowledge acquisition is essential.

The acquisition of knowledge being complete, the representation of that acquired knowledge within the computer becomes the task at hand. The process of knowledge representation is one of translating facts about the world into meaningful symbols with which the computer can work. Three primary means of knowledge representation have been most commonly used in the AI community: rule base (production systems), semantic nets, and frame base representation.

Rule based systems were first prepared in the 1940s and later refined by Newell and Simon. These systems obtain their power by a data base and a set of production rules [Ref. 34:p. 30]. Governed by a control system which determines when the appropriate rule is to be found, the rule based system can both represent knowledge and assert new facts. Knowledge representation in the rule based system is usually in the form of the first order predicate calculus (discussed earlier). The order of firing of the rules may be either forward chained or backward chained. Forward chaining invokes a rule and applies it to factual evidence in the attempt to infer more facts. For example:

> If there is water in the glass and (if) the glass is knocked over, then the water will spill.
>
> Rule + Fact = Conclusion (another fact, goal)

Backward chaining utilizes facts of invoked rules which support conclusions. For example:

> Conclusion: the water is spilled
>
> Fact: the glass has water in it
>
> Fact: the glass was knocked over
>
> Rule: if there is water in the glass and the glass is knocked over, then conclusion
>
> Conclusion + Fact = rule

Thus, in these simple examples, the use of factual evidence and heuristic knowledge is demonstrated, however, the limit of the system is quickly understood if the water in the glass was frozen. A human would recognize this new fact and make the appropriate heuristic adjustment in considering the situation while the ES might miss it completely.

Semantic nets (networks) are a representation of knowledge based on a network structure where nodes, representing facts, are linked by arcs, representing relationships [Ref. 20:p. 70]. For example, a rifle (node) "is a" (arc) firearm (node). These networks can then be manipulated by logical operators in a manner similar to that of the predicate calculus.

The concept of organizing knowledge by frames was proposed by Minsky in 1974 [Ref. 34:p. 39]. A frame may be thought of as a slot which may only be filled by a specific type of knowledge. Knowledge may be entered into a frame, or extracted, or a frame may be left vacant. "As their structure suggests, frame systems are useful for problem domains where expectations about the form and content of data play an important role in problem solving, such as interpretating visual scenes or understanding speech" [Ref. 20:p. 85]. Frame systems may prove to be most useful where a standard format utilizing information specific frames are prevalent, as in report generation.

Facts and rules, facts and arcs, and facts and frames are three generally accepted ways of representing and, thus, making available for manipulation by logical operators knowledge about the problem domain. Also, a structure known as a global data base or a blackboard is generally found in the body of an ES. This construction is used for "keeping track of the problem status, the input data for the particular problem, and the relevant history of what has thus far been done" [Ref. 35:p. 47]. The blackboard works then as a clearing house for communication between elements of the program by supplying needed information to these cooperating entities.

Thus, the ES can be grossly divided into three major parts, the control structure (or inference engine), the method of knowledge representation, and the blackboard. When working correctly, the ES should produce satisfactory solutions based upon its expertise in its problem domain. It is crucial to understand that these solutions are probalistic and not deterministic. The whole field of study of decision making under uncertainty is pertinent and must necessarily to be applied to ES in such a manner that solutions generated resemble as closely as possible those which they are to represent in the real world. But, how will the machine's decisions (solutions) be accepted by decision makers? We have become confident in a computer's ability to calculate quickly and accurately, but will the same confidence be felt as computers take on more important logical processing knowledge? Two ways in which to aid in the acceptance of ES by decision makers are first, assure an agreement by the community of experts with regard to the knowledge represented in the ES and, second, continually validate the decisions of the ES by review, critique, and update of the ES with the assistance of experts in the field. What this leads to is that for all the good news about artificial expertise versus human expertise:

Human Expertise	Artificial Expertise
Perishable	Permanent
Difficult to transfer	Easy to transfer
Difficult to document	Easy to document
Unpredictable	Consistent
Expensive	Affordable

[Ref. 20:p. 12]

There is also bad news about artificial expertise:

Human Expertise	Artificial Expertise
Creative	Uninspired
Adaptive	Needs to be told
Sensory Experience	Symbolic input
Broad focus	Narrow focus

Broad focus	Narrow focus
Common sense knowledge	Technical knowledge

[Ref. 20:p. 14]

And the conclusion:

> For these reasons and others relating to public acceptance of artificial expertise, expert systems are often used in an advisory capacity - as a consultant or aid to either an expert or novice user in some problem area. [Ref. 20:p. 15]

The likelihood of removing the human factor from the decision making loop at this time appears remote.

III. NATURAL LANGUAGES PROCESSING (NLP)

As the title of this chapter suggests, NLP has the ability of a computer to be programmed in a manner capable of understanding human language. If this ability could be achieved by a computer it would open the door to man/machine interaction and essentially remove the boundaries which limit the ability for man and machines to communicate. NLP, if available at the input and output junction of the computer system would eliminate the need for a transitional programming language for interaction with the computer. Unfortunately, NLP may be the greatest challenge confronting AI researches today. Gevarter puts it in the following terms:

> Human communication in natural language is an activity of the whole intellect. AI researchers, in trying to formulate what is required to properly address natural language, find themselves involved in the long term endeavor of having to come to grips with this whole activity. [Ref 19:p. III]

Thus, not only is the immediate advantage of easier and more efficient human machine interface visualized, but our understanding of the activity of the whole human intellect may be expanded.

Once NLP comes into full development, its applications are numerous:

-Speech Understanding

-Story Understanding

-Information Retrieval

-Question Answering Systems

-Computer Aided Instruction (CAI)

-Machine Translation

-Document and Text Understanding

-Automatic Paraphrasing

-Knowledge Compilation

-Expert Systems Interface

-Decision Support System

-Explanation Modules for Computer Actions

-Interactive Interface to Computer Programs

-Control of Complex Machines

-Document or Text Generation

-Speech Outputs

-Writing Aids

But, how is the problem of the enormous task of programming a model to simulate human understanding of language to be built?

The original attempts at NLP were via a mechanical dictionary or machine translator. A. Oettinger is attributed with creating the first mechanical dictionary [Ref. 6:p. 3]. The early hope was that the mechanical translator would permit the computer to undertake direct translation of foreign languages. However, this problem proved more difficult to solve than first thought.

> One possibility that any scientist must take into account is that he or she has made a poor choice of problem. A poor scientific problem is one that cannot be solved with the knowledge and tools available at the time. The "classic" example of this is the alchemists' attempt to change lead into gold. A more recent example,--is the attempt to do automatic translation in the late 1950's and early 1960's. [Ref. 9:p. 172]

And, machine translation proved to elude the AI alchemist as much as gold had the traditional alchemists in times past. The matter of machine translation seemed a simple one given the power of the digital computer, as for every word in one language there was to be a similar one in another language such that direct translation would take place. And, what if words became ambiguous? Well, simply print all possibilities of translation. This method did not solve the problem of dealing with ambiguity of words. "Instead, researchers started working on phrase-by-phrase , and sentence-by-sentence translations. . .[Ref. 9:p. 172].

There were some modest successes in the ares of NLP as demonstrated by a program developed in 1954 at Georgetown University. The program made use of six syntactic rules and a 250 word dictionary. Its meagre success was quickly overrated and many other experiments in the NLP sprung up at Harvard, MIT, and the University of Pennsylvania. The goal was to conduct direct translation which would be proofread and put in a finished form by a human being. The interest in this work was so keen that five government agencies spent some $20 million on this area of research by 1966. As early as 1963, a government report, "Languages and Machines" concluded:

> We have already noted that, while we have machine aided translation of general scientific text, we do not have useful machine translation. Furthermore, there is no immediate or predictable prospect of useful machine translation [Ref. 6:p. 4]

A critique by the influential linguist, Bar-Hillel, in 1964, pointed out that there was. . ."no way to do word sense disambiguation without deep understanding of what the sentence meant" [Ref. 9:p. 173]. In 1966, at the request of the National Science Foundation, the Pierce Report (Pierce, 1966) was published which concluded that, "There was no way that the work on machine translation could be justified in terms of practical output" [Ref. 9:p. 173]. The products of such translations were of poorer quality and more costly than those produced by human translators.

This lack of success lead Gervarter to flatly state that "By 1970, AI had only limited success. Natural Language Translating had already collapsed" [Ref. 19:p. 10]

Perhaps the most classic and illustrative example of this collapse follows:

> The effort (Language translator) was a failure. When the sentence "The spirit was willing but the flesh was weak", was translated into Russian and back into English,

> it is said to come out as "the Vodka is strong but the meat is rotten. [Ref. 19:p. 130]

Work then on mechanical translators was mostly abandond in the 1960's; research, however, was continued in natural language understanding--research "directed toward the automatic comprehension of the English language in which people habitually think and communicate." [Ref. 26:p. 13]. The keys to this understanding were provided by early researchers, that is to understand language not word by word, but phrase by phrase and sentence by sentence. These keys are lessons which come slowly, with difficulty, at great momentary expense, and those in AI seem to have difficulty remembering this. As Charniak and McDermott point put:

> It is sobering to consider the difference in quality between experimental translations done in 1956 and the "real" ones done nearly ten years later. The experimental versions are much clearer. The reason, of course, is the very limited domain, vocabulary, and syntax used in the experiment. The simple ideas do not necessarily scale up to real world problems, and this is a lesson we in Artificial Intelligence have been taught many times. The trick is to remember the lesson each time before investing too heavily in a particular bag of ideas. [Ref. 9:p. 174]

Thus, the solution to NLP, as with all areas of AI, is not a simple one of direct word to word translation but compounds exponentially as one examines the nature of the problem. For, as was stated in the beginning of this chapter, NLP of human communication is an attempt to understand an activity of the whole intellect. In searching for a method of solution to NLP, the problem, as others in AI, must be subjected to dissection into smaller parts. The following discussion will not attempt to offer solutions to the problem of NLP or its sub-problems, but will describe the search areas and difficulties associated with each. Much of this description will come via definitions.

Attempting to understand language on a word by word basis is like attempting to describe the proverbial forest by a description of each tree. Thus, there must be a mechanism for sorting out areas of trees which have descriptive meaning in context of the forest or in this case, sentence structure which has meaning in context of the sentence as a whole. The vehicle for obtaining these sentence structures is a parser and the structures themselves are the syntax of the language.

A parser is generally intended as a formalism that assigns a structural description to a sentence; also used to describe formalisms that assign a semantic interpretation to a sentence (or a parser for a semantic grammer) [Ref. 36:p. 269]. Or, perhaps a more easily understood definition "1. to break (a sentence) down into parts, explaining the grammatical form, function, and interrelation of each part. 2. to describe the form, part of speech, and function of (a word) in a sentence" [Ref. 30:p. 1065].

Thus, by parsing a sentence, more managable, and hopefully, more easily understandable sentence structures are provided to a NLP program for manipulation. Syntax is simply "the structural description of a language" [Ref. 36:p. 270].

A parser must accomplish two objectives: it must upon dissecting a sentence, determine if it meets the grammetical requirements of the language; it must then create a representation of this grammetical structure. This representation is most often in the form of an augmented transition network, (ATN). Such a network can be used as an internal representation for the computer and, when illustrated, become a graphical notation consisting of nodes and arcs. A node indicates the grammetical (syntaxtic) structure parsed and an arc represent the logical connections between two nodes. Thus, sentence structures such as subject (S), noun phrase (NP), adjective (A), verb (V), verb phrase (VP) may be mapped from a parsed sentence. Parsing and its graphical representation of any augmented transition network may be, in a simplistic way, viewed as the skeletal structure of a diagrammed sentence. The first job of the NLP is to parse the sentence, construct an ANT and fill in the nodes (which contain the correct syntactic construction), and associate the node via arcs.

There are two basic methods of parsing: top-down, and bottom up. An ATN which makes use of top-down parser by "making implicit expectations of what will be found next in the sentence, based on what has been found...ATN grammars can be implemented so that the most probable (or least expensive) choices are considered first, thus minimizing backup. This is called heuristic parsing [Ref. 31:pp. 66-67]. Top-down parsers have several advantages and disadvantages. Their advantages are that "they are easy to write, and they can be ordered heuristically. . ." [Ref. 36:p. 72]. Their disadvantages are that the same syntactic structure may be parsed several times, even though correctly parsed the first time, as the parser backtracks to complete operation on the sentence. Secondly, a top-down parse must work on a sentence which is clearly, bounded, as a successful parse is achieved only when the end of a sentence is reached. Third, if the parser fails, then backtracking can be exhaustive and non-indicative of where the road block to parsing the sentence lies. Fourth, if a structure in the middle of the sentence cannot be parsed, the remainder of the sentence will remain undertermined [Ref. 36:p. 75]

Bottom-up parsers are data driven devices. These parsers follow an input string from left to right building all possible syntactical structures to the left of the pointer as the pointer moves to the right, word by word. Thus, the bottom-up parser is data driven and prone to develop sub-structures of the sentence which may never be used. However, this type of parser allows for ambiguities in the sentence. The main drawback of a bottom-up parcer is that "if a new sentence must be written for each version of each phrase, the number of combinations for long sentences grows out of hand quickly" [Ref. 36:p. 76].

The primary trade-off between the top-down and the bottom-up parsers is one of efficiency. However, by accepting the efficiency of the top-down method, one gives up the flexibility and completeness of the bottom-down method.

The preceeding explanation will provide a rough idea of how sentence structure plays a role in understanding language and NLP programming. From here it is important to gain an appreciation of the symbols of language or words.

Barr and Feigenbaum (1981, p. 332) define semantics as "the meaning of words and sentences [Ref. 19:p. 115]. However, this simplistic definition may be very misleading:

> Semantic processing (as it tries to interpret phrases and sentences) attaches meanings to the words. Unfortunately, English does not make this as simple as looking up the word in the dictionary, but provides many difficulties which require context and other knowledge to resolve [Ref. 19:p. 115].

This context and other knowledge are addressed as multiple word senses, pronouns, ellipsis, and substitution, each of which will be discussed separately.

Semantics is the relationship between symbols and concepts. In normal conversation, humans deal quite easily with the meaning of words where computers fail rather quickly at the task. In the case of multiple word senses one need not look far to find examples of dual meaning for the same word (a synonym) e.g., bar, a drinking establishment, and bar, as a piece of metal, wood, etc., which is longer than it is wide. The ambiguity is complicated even more as the word takes on not only a different meaning but a different form of speech as with the verb bar, to block or obstruct. Tennant offers a list of interesting words to consider supplied to him by a friend, Gene Lewis:

> . . .dog, cow, badger, squirrel, fly, horse, buffalo, chicken, and snake. [Ref. 36:p. 103]

A few moments of reflection and the reader will realize the complications presented by this list as they take on different meanings and different forms of speech. Humans deal rather easily with these changes by reviewing the word in question in context. Not only does context allow humans to understand the meaning and use of a word with which one is familiar, but also assists one in comprehending the meaning of words whose definitions are unknown. The familiar response when one is asked the meaning of a word which is unknown, "Please place it in context". However, the computer is not afforded the luxury of any form of intellectual intercourse. Early NLP circumvented the problem of multiple meanings and uses "by restricting the domain of discourse so severly that it was highly unlikely that a word would be used in more than one way" [Ref. 36:p. 106]. It is the suspicions of these writers that the problem of multiple senses and uses are treated in a similar fashion by many of today's programs. However, there are conventions to deal with this problem which generally assign primary, secondary, etc., meanings to a word and carry these multiple meanings along until a contextual reference provides some assistance in choosing the most correct word meaning. A means of the choosing may be as an associative construct which will allow the notion of buck (as in dollar bill) to be associated with the word pocket where it would exclude from consideration the notion of buck (as in male deer) to be associated with the word pocket [Ref. 36:p. 168]. One will quickly recognize the complication of association as the size of the computer dictionary grows and the operational domain expands.

The use of pronouns in language is not only frequent but provides for simplified cohesion of text. Without their use, text seems awkward and clumsy. Their use "allows a simplified reference to previously used (or implied) nouns, sets or events" [Ref. 19:p. 115]. Thus, the use of pronouns aid humans in the way they write and speak. For example:

> Joe stopped by Bill's place on his way home. While there he had a beer.

Easy enough for humans to understand but how about the computer? How would it choose the referent of he, Joe or Bill? Thus, the computer might prefer the following version:

> Joe stopped by Bill's place on Joe's way home. While at Bill's place, Joe had a beer.

Although this sentence may appear more thorough in its handling of Joe and Bill, a new confusion may set in as to whether Joe and Bill in the first sentence are the persons as in the second sentence. These examples point to simple cases of reference, and what is the programming convention to handle even these simple cases?

> Pronoun handling is a difficult problem in natural language processing even for the least exotic occurances. Most of these more difficult cases are beyond the capabilities of current systems, but work is proceeding along these lines. [Ref. 38:p. 117]

Ellipses and substitutions represent two other language conventions which must be dealt with by a NLP. An "ellipsis is the phenomenon of not stating explicitly some words in a sentence, but leaving it to the reader or listener to fill them in. Substitution is similar--using a dummy word in place of the omitted words" [Ref. 19:p. 115].

For example:

> Joe caught three flies. Bill caught two. (elipsis)
> Joe uses a right handed glove, Bill uses a left handed one. (substitution)

Though ellipsis and substitution present their unique problems to NLP, they can, "By Employing Pragmatics, ellipsis and substitutions can usually be resolved by matching the incomplete statement to the structure of previous recent sentences, finding the best partial match and the filling in the rest from this matching previous structure" [Ref. 19:p. 115]. No simple task and one, again, which requires considerable memory and iteration before the "best " match is made.

The final topic in the cursory view of NLP is pragmatics, "the study of the role of contextual knowledge in language; knowledge about the world" [Ref. 36:p. 269]. Obviously a knowledge--dependent subject which can consume every bit of memory available and offer enormous problems in indexing, cross-referencing , and utilization. Most humans merely scratch the surface of pragmatic knowledge and where some depth is achieved it is through specifying and limiting oneself to a particular domain. Thus, the only way to treat the difficulty of pragmatic knowledge in a computer is in a similar manner and that being the impositon of boundaries on both knowledge and problem domains.

This concludes a brief review of NLP but the reader is warned that only the top layer of the subject has been addressed in order to give a flavor of the nature of NLP. Other areas which need to be explored by researchers in NLP are listed:

Morphology

Phrase Structure Grammar

Context Free Grammar

Transformational Grammar

Case Grammar

Semantic Grammars

Pragmatic Ellipsis

Structural Ellipsis

Procedural Representations

Declaration Representations

Case Frames

Conceptual Dependency

Frame

Script

Speech Recognition

Speech Understanding

Modifier Attachment

Noun-Noun Modification

Decomposition

Speech Acts

Rules of Dialogue

The list continues. To produce a NLP which aids in the interaction between humans and machines is probably within our grasp, and many efforts are in place today. However, to achieve a model of natural language intercourse between humans and machines is as remote as our complete understanding of the cognitive processes of humans, the obstacle which hinders all AI efforts.

IV. COMPUTER VISION AND PATTERN RECOGNITION

It is important to preface this chapter with a paragraph regarding its content. The preceding chapters have been highly descriptive in nature in order to serve the purpose of informing the reader of the technologies in AI and, to an extent, how they may be useful considering their strengths and weaknesses. The content of this chapter can become much more technical much more quickly. To avoid becoming involved in the technicalities of the "how it is done", this chapter like the others, is highly descriptive. Thus, the reader whose appetite is whetted by the subject is encouraged to look more deeply into visual and pattern precessing via several mathematical models including Fourier Transforms and Gaussian statistics. Two good books with which to start this continued exploration are Pattern Recognition by Satoshi Watonabe and Artificial Intelligence by Charniak and McDermott.

Pattern recognition and computer vision may, at a high level, be defined separately. Pattern recognition is as follows:

> Be it a blob or lines or something without a name, a pattern is the opposite of chaos; it is an entity vaguely defined, that could be given a name. . .i.e., a something. . .[Ref. 40:p. 2]

Essential in this definition is the notion of something, not chaos. Thus, one should not limit his mental image of pattern to only those generated visually, but extend the notion of pattern to include those produced by electromagnetic emmissions, vibrations, odors; IR signals, etc.

Computer vision which may be thought of as a subcategory of pattern recognition is defined by Barrow and Tenenbaum (1981, p. 573).

> Vision is an information processing task with brightness values, representing projections of three-dimensional scenes recorded by a camera or comparable imaging device.

[Ref. 35:p. 84]

Therefore, in many ways the principles of explanation which apply to the one apply equally to the other topic. Where this is not the case, it will be noted.

There are three primary models from which to view pattern recognition; paradigm matching, associative recognition, and constructive recognition [Ref. 39:Ch. 1].

Given a general description of a pattern i.e., a model or paradigm, and upon observing an object which is similar to the model in its features one may conclude that it fits into the category of the model. For example, if one is given the definition of a square as a closed curve, with four equal length straight sides, intersecting at 90^{o} angles, and upon discovering a phenomenon which approximates these criteria, may conclude that such a phenomenon is indeed an object of the class 'square'.

Pattern recognition by association is quite simple. Given a person who has never seen a tree and is placed in front of a forest he may quickly recognize that many of the entities in the forest have properties similar to those of trees. From observation, one may then conclude that those entities with similar properties may be associated, and generalized, and conclude that the entity represents the class of object, the tree.

Finally, patterns may be generated and recognized by an enumeration of properties until the collection of these properties construct a recognizable pattern. For example, it has four sides of equal length, these sides intersect at 90^0 angles, it is a closed curve. Thus, a good guess is that it is a square.

It is interesting to note that each model just described may have a counterpart in the notions of knowledge representation in expert systems:

paridigm matching-frames

constructive recognition - rule based

associative recognition - semantic networks

The reader is encouraged to try to determine if the constructive and paradigm models are not one and the same.

Using any of these models requires data, for pattern recognition and computer "vision is an information processing task" [Ref. 40: p. 19]. The collection of data is the first step in pattern recognition. Since scenes of patterns are generally too complicated to interpret as a whole they must be subdivided into small cells of data or pixels (picture elements) in the case of vision. The determination of how small these cells must be is a function of the pattern to be recognized. It may not be necessary to describe each individual straw in order to recognize a haystack. But, each cell must contain enough data to contribute to recognizing the pattern without obscuring it. The case of one being unable to see the tree for the forest.

Finally, the data in the cell must be expressable in a binary system if it is to be useful for computer analysis. An important point regarding computer processing is that as the number of cells for processing increases so does processing time. And, as the number of cells become very large, serial processing quickly becomes overwhelmed. It has been suggested that human vision must utilize a form of parallel processing, and if computer pattern recognition is to become robust it, too, must use parallel processing [Ref. 35: p. 42].

Once the data has been collected it must be submitted to two levels of processing, early and late:

> In early processing, the goal is to get useful information from the raw data image, and every part of the image is

> processed in the same way. In late processing, the goal is to find the objects from the useful information. [Ref. 9:p. 95]

Early processing must then supply to the constructive model, previously discussed, data which provides for the recognition of an element. Late processing data (elements) may be most useful to the associative and paradigm models. In fact, early and late processing are continuing, complimentary functions which allow a scene to be sketched. This is particularly true for computer vision.

If a computer is to be successful in replicating human vision it must be provided with information to create virtual lines, interpret texture and shading, and conduct motion and stereo analysis. Of these types of information, the construction of virtual lines is fundamental as these will construct the primal sketch of the image, a wire-line drawing or dot-to-dot drawing. Then this drawing may be enriched with the information available from the other categories.

The process described requires the computer to both differentiate and integrate in successive steps. As the pieces of the known whole begins to emerge then they must be then reconstructed in order to gain a meaningful interpretation of the whole. If some piece of information is missing the computer may find that the process of integration may be very frustrating. The human being on the other hand, deals with incomplete integration by anticipation or expectation. In instances where complete information was available complete integration takes place and later, under similar circumstances, when incomplete information exists, complete integration could still take place based upon expectation.

The process that both humans and computers undertake in pattern recognition is to collect specific data from cells (differentiation), create elements from these cells, and create patterns (integrate) from these elements. Thus, factoring the whole, filtering the data from noise, limiting the amount of data so as not to overwhelm the system, and grouping the data play intricate and vital roles in pattern recognition.

This description of pattern recognition for all purposes supposed a static rather than a dynamic scene. Add to the problem of pattern recognition the elements of time and motion and the difficulty of the problem grows astronomically. And although, the principles of pattern recognition and vision may appear simple at first glance, in-depth understanding may be lacking. Charniak and McDermott cite Barrow:

> Despite considerable progress in recent years, our understanding of the principles underlying visual perception remains primitive. Attempts to construct computer models for the interpretation of arbitrary scenes have resulted in such poor performance, limited range of abilities, and inflexibility that, were it not

> for the human existence proof, we might have tempted long ago to conclude that high performance, general-purpose vision is impossible.

"We complete this dire observation by remarking that most of the results...are no more then ten years old. Although they look quite promising they are fragmentary and not secure" [Ref. 9:pp. 94-95].

As with the other categories of AI, pattern recognition and computer vision may be more promise than substance. But, even the promise holds hope for application in the areas of robotics, inspection tasks, remote sensing, tracking moving objects, navigation (passive), aid to the partially sighted, and etc. [Ref. 9:pp. 94-95].

Finally, the same question must be asked of computer vision and pattern recognition as is asked of all areas of AI: when will it live up to its promise? The answer must be measured in evolutionary steps and not evolutionary leaps. This means that if one is awaiting computer vision (or its comparable counterpart in ES, NLP, or Robotics) to duplicate human vision, the wait may be a long one.

V. ROBOTICS AND AUTONOMOUS VEHICLES

It should come as no surprise that robotics as with other elements of science is an evolutionary entity from the simple tool to complex machines to automated machinery. Consequently, the line separating robots from autonomous vehicles becomes blurred. The Robot Institute of America defines a robot as a "reprogrammable, multifunctional, manipulator, designed to move material, parts, tools, or specialized devices, through variable programmed motions for a variety of tasks" [Ref. 35:p. 159]. However, this definition must be broadened in scope when AI is introduced into the system so that the word 'programmed' is not meant simply as preprogrammed actions which the robot undertakes. For with the introduction of AI techniques to robotics the system may then be able to respond to its environment by planning, executing, and evaluating its actions. Thus, the autonomous vehicle is a device which can essentially solve problems associated with a task, from start to satisfactory completion, without external contributions provided by other machines or human beings. Those devices which are influenced or controlled by external inputs are called teleoperated or telepresence devices. These are "machines capable of action at a distance under the control of a human being". [Ref. 35:p. 8], [Ref. 42:p. 170]

A robot is generally equipped with the following sub-systems:

- one or more manipulators (arms)
- and effectors (hands)
- a controller
- and, increasingly, sensors to provide information about the environment and feedback of performance of task accomplishment [Ref. 35:p. 159].

The manipulators may work in two or three dimensional space presented by cartesian, cylindrical, polar, articulated or joint spherical (revolute) coordinate systems. Each configuration provides the prescribed range of motion necessary to perform a particular task. As the task becomes more complex and a greater range of motion is required, the more complex the coordinate system becomes, requiring greater computational strength [Ref. 35:p. 168]. Execution of motion in any of the coordinate systems may be accomplished by electrical, pneumatic, or hydraulic means.

End effectors are essentially the hand of the robot which allows it to undertake tasks. The greatest difficulty with these effectors is the lack of a wide range of motions.

Control of the robot may be undertaken by something as simple as a system of mechanical devices with prescribed settings and stops which are an integral part of their tooling, or by computer programs which control the robot's actions. As complexity of the task increases and greater task flexibility is required, the more sophisticated the control device must be.

Sensors provide the robot a means of interacting with its environment. These sensors may include sight, touch, temperature, smell and hearing, all of which are at various stages of development. Something as simple as a bi-metallic switch may act as an adequate sensor of temperature, offering simplicity and reliability. A sensor for sight, which is critical for many applications of robots e.g., object recognition and avoidance, may be very difficult to construct.

> How soon before we get human-like vision is hard to say. If this intelligent system of vision is, as we believe, built up from the 200 modules of a visual task, it will take 10 doctoral theses to explain and unravel each one. In terms of research work, this means about 2,000 of the right theses published a year, of which probably half about vision. So, given our current rate of progress, it will be about 20 years before a truly human vision system is realized in a machine. [Ref. 43:p. 118]

Given the use of a camera it provides its own difficulties in operation.

> For instance, our present software for the visual obstacle course has a camera calibration phase in which the robot parks itself precisely in front of an exact grid of spots so that a program can determine a function that corrects for distortions in its camera optics. This allows other programs to make precise measurements of visual angles in spite of distortions in the camera lens. We have noticed that our present code is very sensitive to miscalibrations, and we're working on a method to continuously calibrate the camera from the images perceived on normal trips through clutter. With

> such a procedure in place, a bump that slightly shifts one of the robot's camera will no longer cause systematic errors in its navigation. Animals seem to tune most of their nervous systems with processes of this kind, and such accomodation may be a precursor to more general kinds of learning. [Ref. 44:p. 133]

Research is currently being conducted by the Automation Technology Branch (ATB) of NASA, Langley Research Center involving use of "telecoperator (remote controlled manipulation), robotic (autonomously controlled manipulation), devices for remote space application...In order to realize this technology, the ATB is conducting research in six major areas:

(1) manipulator dynamics and control

(2) end effectors

(3) sensors

(4) operator-machine interface to automated system

(5) distributed computer and network systems

(6) Artificial Intelligence

[Ref. 45:p. 1]

This list of technological development is similar to Gevarters; however, it places its emphasis upon the future of robotics which is autonomy through advances in computational power, and AI which will integrate perception, reasoning, and manipulation. Harrison and Orlando point out how the autonomous device differs from the traditional robot.

> The successful implementation of a manipulator system requires that the system contains elements that can simultaneously and accurately perceive, reason about, and interact with its environment. In the traditional approach to teleoperation, the perception and reasoning elements are the human operator, who is totally dedicated to the execution of a task. On the other hand, the traditional robotic system is performance limited by its lack of perception and reasoning capabilities. A logical alternative is the development of teleoperator devices that can accomplish operator selected task primitives and that can be expanded over time through intelligent automation techniques toward the realization of relatively autonomous, or robotic operation [Ref. 46:pp. 8-9].

In another of her papers Orlando states the necessary requirements to achieve autonomous operations which are knowledge representation, intelligent planning, operator-machine interface, and system integration [Ref. 47:p. 2]. Perhaps the most important of these requirements is systems integration. For underlying the notion of autonomous device is the concept of cooperating subsystems communicating with each other via a blackboard.

> . . .planner, perception, and control subsystems are regarded as a community of cooperating entities. These entities are coordinated through the exchange, for plans and reports. Intelligent communication interfaces (ICI) for each module use reports to maintain the consistency of local copies of a distributed blackboard. In this structure plans can be exchanged between any of the modules...[Ref. 48:p. 17]

Orlando discusses in detail and suggests the software implementations necessary to achieve autonomy but leaves open the question as to what degree of machine intelligence is necessary for the successful completion of a particular task. . .[Ref 46:p. 9], [Ref. 45], [Ref. 47]

Hans Moravec in his article The Rovers offers his critique of the state of the art of sensory and control systems and speculates upon the necessary computational power which may be necessary to make an autonomous system operational.

> While our sensory and muscle-control systems have been in development for a billion years and common sense reasoning has been honed for probably about a million, really high-level, deep thinking is little more than a parlor trick, culturally developed just a few thousand years ago. A few humans, operating largely against their natures, can learn this trick. As with Samuel Johnson's dancing dog, what is amazing is not how well it is done, but that it is done at all.
>
> Computers can challenge humans in intellectual areas where humans perform inefficiently, because they can be programmed to carry on much less wastefully. An extreme example is arithmetic, a function learned by humans with great difficulty, but which is instinctive to computers. These days an average computer can add a million large numbers in a second, which is more than a million times faster than a person, and with no errors. (And yet, one hundred-millionth of the neurons in a human brain, if reorganized into an adder using switching-logic-design principles, could sum a thousand numbers per second. If the whole brain were organized this way, it could do sums one hundred thousand times faster than the computer!)
>
> Computers do not challenge humans in perceptual and control areas because these billion-year-old functions are carried out by large portions of the nervous system operating as efficiently as the hypothetical neuron adder above. Present-day computers, however efficiently programmed, are simply too puny to keep up. Evidence comes from the most extensive piece of reverse engineering yet done on the vertebrate brain, the functional decoding of some the visual system by David H. Hubel, Torsten N. Weisel, and their colleagues at MIT.
>
> The vertebrate retina's 20 million neurons take signals from a million light sensors and combine them in a series of simple operations to detect things like edges, curvature, and motion. The image thus processed

is sent to the much bigger visual cortex in the brain. Assuming the visual cortex does as much computing for its size as the retina, we can estimate the total capability of the system. The optic nerve has 1 million signal carrying fibers, and the optical cortex is a thousand times thicker than the layer of neurons that do the basic retinal operations. The eye can process 10 images/second, so the cortex handles the equivalent of 10,000 simple retinal operations a second, or 3 million-/hour.

An efficient program running on a typical computer can do the equivalent work of retinal operation in about the two minutes, for a rate of 30 per hour. Thus, seeing programs on present-day computers seem to be 100,000 times slower than vertebrate vision. The whole brain is about 10 times larger than the visual system, so it should be possible to write real-time human equivalent programs for a machine 1 million times more powerful than today's medium-size computer. Even today's largest super-computers are about 1,000 times slower than this diseratum. How long before our research medium is rich enough for full intelligence?

Since the 1950s, computers have gained a factor of 1,000 in speed-per-constant-dollar every decade. There are enough developments in the technological pipeline to continue this pace for the foreseeable future.

The processing power available to AI programs has not increased proportionately. Budget increases spent on convenience features--operating systems, time-sharing, high-level languages, compilers, graphics, editors, mail systems, net-working, personal machines, etc.--have been spread more thinly over even greater numbers of users. I believe this hiatus in the growth of processing power explains the disappointing pace of the development of AI in the past fifteen years, nevertheless it represents a good investment. Basic computing facilities are now widely available, and thanks largely to the initiative of the instigators of Japanese Super-computer and Fifth Generation Computer projects - attention world-wide is focusing on the problem of processing power of AI.

The new interest in computing power should ensure that AI programs share in the thousand fold-per-decade increase from now on. This puts the time for human equivalence at twenty years. Since the smallest vertebrates, like shrews and hummingsbirds, produce interesting behavior with nervous systems one ten thousandth the size of

> human's, we can expect fair motor and perceptual competence in less than a decade. By my calculation and impressions, present robot programs are now similar in power to the control systems of insects.
>
> Some principles in the Fifth Generation Project have been quoted as planning "man-capable" systems in ten years. I believe this more optimistic projection is unlikely, but not impossible. The fastest present and nascent computers, notably the super-computers Cray X-MP and Cray 2, compute at 10^9 operation/second, only they do it 1,000 times too slowly. [Ref. 45:p. 133-136]

Obviously, as Moravec points out, an enormous undertaking, so why seek to build an autonomous vehicle or robot? First, an autonomous vehicle will provide relief for humans from monotonous tasks. Second, it will replace humans in dangerous environments. Third, it will provide graceful degradation of its mission over time, unlike the teleoperated system which becomes dysfunctional when its operator-machine communication links are severed. Finally, it may provide some economic relief: "The Apollo project put people on the moon for $40 billion. Viking landed machines on Mars for $1 billion". [Ref. 44:p. 188]

VI. AN INTRODUCTION TO SPACE SYSTEMS

The following few pages are a brief description of a generic model of a space system. The description will be divided into three sections: vehicular systems, mission systems, and ground node systems and considerations. The purpose is to offer an overview of the major component issues which are inherent in designing and building any space system. Specifics are deliberately avoided for two reasons, first the highly technical nature of engineering design, which is beyond the scope of this paper, and secondly, where greater specification might prove helpful it usually becomes system specific. Even in a series of four or five spacecraft, unique changes will be introduced to each platform. There is no mass assembly line for the production of satellites today in the U.S. The following description relies very heavily upon the class notes written by Distinguished Professor Allen Fuhs and used in his upper level graduate course in Aeronautical Engineering, Space Craft Design, which he teaches at the Naval Post-graduate School.

A. VEHICLE SYSTEMS

A space system is most easily viewed as a group of subsystems, components and elements. The subsystems may be generally considered the launch facility, (a cogent discussion of launch facilities and satellite orbitology is presented in the thesis, "Military Applications of Space: An Introductory Text", written by Beth E. Patridge, Lt. USN, Naval Postgraduate School, June 1985), the launch vehicle, and the launch bus which provides the platform for the mission payload.

B. LAUNCH VEHICLES

There are a wide variety of launch vehicles available today since their production development began in the 1950's. Several features have become desirable, and are usually incorporated, in these solid rocket boosters:

1) high energy release per unit mass
2) high density
3) low pressure
4) low sensitivity of burning rate to temperature
5) flexibility in controlling burning rate
6) ease of ignition
7) reproducibility of performance
8) good stability in storage
9) resistance to detonation
10) low toxicity of exhaust products
11) ease of processing and handling
12) raw material readily available at low cost

[Ref. 49:p. 6]

most of the above characteristics are met by today's rockets. But, the single overriding expectation which the launch vehicle must meet is the ability to place a payload safely and accurately into space, reliability. Today's reliability stands in the 90% - 98% range [Ref. 50:p. 204]. So good are these rockets that Vogel concludes "perhaps it is a fitting tribute to the people who design and build solid rockets that reliable operation of their products are virtually taken for granted" [Ref. 49:p. 29].

C. SPACE CRAFT BUS

The space craft bus is the platform which houses the mission payload and performs a variety of "house keeping" functions. Each function will be addressed individually.

D. ALTITUDE CONTROL

Once a space craft is placed into orbit that orbit immediately begins to decay due to aerodynamic drag, gravity and other environmental factors. Aerodynamic drag and gravity are the principle cause of orbital decay, and fall off inversely to the square of the spacecraft is altitude. This decay must be compensated for by thrust motors fired for specific times to maintain the desired orbit. Although the thrust motors are an integral part of the space craft the determination of when and for what duration to fire these motors is a function of the ground control station.

E. ATTITUDE CONTROL

Once the decay of a satellite orbit has been successfully contended with, the attitude at which the craft is flying must be kept stable, or adjusted, as mission requirements dictate. In order to achieve this attitude control a number of different techniques may be employed. These techniques include mass expulsion (pneumatic systems), momentum storage devices, gravity, gradient systems, spin stabilization, and magnet systems. Each technique has advantages and disadvantages, but all serve the same purpose: keeping the space craft properly oriented in orbit.

1. Spin Stabilization

By spinning a spacecraft about its axis of maximum inertia in, the absence of applied torque, provides stabilization similar to that of a gyroscope, and provides fixed inertial orientation with limited accuracy for negligible weight. This type of stabilization may be inappropriate for optical line of sight requirements

2. Magnetic Systems

Stabilization of a spacecraft may be achieved by producing magnetic fields in loops, on board the craft which align themselves with the earth's magnetic field.

3. Gravity Gradient

Gravity gradient control is a simple, passive, and reliable means of attitude control. The principle upon which it works is:

> The difference in the earth's gravity field at the top and botton ends of the space vehicle creates a tongue which aligns (the) vehicle with the local vertical. A damper is used to reduce oscilations. [Ref. 51:p. 2]

Although simple, the reliable gravity gradient control is extremely sensitive to environmental torques and payload motion.

4. Mass Expulsion

Mass expulsion systems are usually pneumatic, utilizing the expulsion of gas under pressure through control jets in a closed loop system. This type of attitude control is insensitive to disturbance torques and provides the widest variety of control orientations. The heavy weight of the system is its greatest disadvantage, particularly for missions of long duration.

5. Momentum Storage Device

These devices usually take the form of power driven reaction wheels, gyros, or fluid filled loops which may absorb disturbing torques or impart correcting torques into the spacecraft. These systems have no expendable fuel requirement, have very precise nulling control, and their precision is only limited by their attitude or sighting sensors. However, momentum storage devices require a means to unload momentum. This produces a highly reliable system with minimum weight and no requirement for attitude sensors. Its accuracy is limited to a few degrees and use is limited to altitudes below 20,000 nautical miles.

F. THERMAL CONTROL

As the spacecraft passes in and out of the shadow of the earth and its attitude toward the sun changes, temperature will fluctuate. Also contributing to this fluctuation is the heat dissipated by power consumption of the internal operations of the spacecraft's equipment. The result of a radical temperature change can be equipment failure due to the temperature being out of equipment operating range. Failure may take place in electrical or electronic components, fuel lines, or other components of operating systems. To prevent equipment failures, temperature must be controlled by the use of shielding radiators and insulation.

Shielding louvers may be heat absorbent or heat reflective or a combination thereof. As the spacecraft becomes hot or cold from changes in its orientation to the sun, these louvers or shields may be positioned to provide heat or proteccion from heat, as necessary. Heat generated internally may be radiated out of the spacecraft or channeled to other parts of the space vehicle where additional heating is required.

G. SURVIVABILITY

Generally speaking, there are two groups of threats that a spacecraft must survive, manmade and environmental. Manmade threats are either intentional or unintentional. Unintentional threats consist primarily of space junk which might damage or destroy the spacecraft. To avoid this type of threat the space craft should be placed in a junk free environment or provided with adequate maneuvering capability to avoid random collisions with large masses. Surviving intentional man-made threats is more difficult for as the threats increase (laser, nuclear burst, charged particle beam) so the defenses must also increase.

Environmentally, spacecraft must primarily survive natural radiation and micro-meteoroid showers. To protect against radiation, the electrical and electronic components must be provided with shielding. To protect against micro-meteroiod hits, the shell of the spacecraft must survive the gradual deterioration due to millions of strikes. Construction materials, thickness of the the materials and construction of the spacecraft's outer wall play a vital role in creating a survivable spacecraft.

H. POWER SYSTEM

The last major component of the spacecraft bus which will be addressed is the power system. Power systems include nuclear dynamic, chemical dynamic, cryogenic chemical dynamic, fuel cell, solar dynamic, solar static battery and photovoltaic. Most commonly, a spacecraft power system will consist of a combination of battery power and photovoltaic power. The requirements placed upon this system are easily understood. The batteries provide power for the spacecraft during periods when photovoltaic cells are not radiated by sun light. When the solar panels are exposed to the sun they produce electricity for hotel use as well as for storage. The duty cycle of the spacecraft's orbit governs when power is stored in the batteries for later use. Once production and storage of power as electrical energy is controlled within acceptable bounds then the issue of power management or budgeting becomes a main concern. What this amounts to is that there are more requirements for power aboard a spacecraft than can be provided for simultaneously. Therefore, the power must be budgeted and distributed to subsystems on a priority basis.

I. MISSION PAYLOAD

1. Sensors and event detection

The devices for sensing and event detection on board a spacecraft are limited only by what the physics of the situation dictate; therefore, one would expect to find infrared sensors, radar, lasers, and cameras. The type of sensor utilized is specified by the mission of the satellite. In some cases, such as environmental/atmospheric research satellite, multiple sensors may be employed. Whether single or multiple sensors are used, virturially all of the same considerations are in effect, namely power consumption, sensor priority utilization, and data management. Regardless of the sensor, its output must be managed. The data produced will either be processed on the satellite, stored and dumped, or transmitted directly to earth station. Any of these functions are computationally intensive, which means an additional drain on power resources available. While the direct transmission of data to an earth station may be most efficient when viewed from a power standpoint, it may not always be possible--as when the earth station is out of the field of view of the satellite. It may prove advantageous to provide for some on board processing of data so that extraneous or useless data is eliminated at the space node before being transmitted back to earth. On board processing may also be useful in survivability by recognizing and avoiding potential hazards.

2. Weapons

Weapons aboard satellites have not been of general concern historically but with the advent of antisatellite systems and the Strategic Defense Initiative, weapons both defensive and offensive may take on a more prominent role in the future development of spacecraft. The addition of weapons to any space platform entails additional weight, on board processing of their effective use, greater demands upon the spacecraft power systems and greater system cost.

J. GROUND SYSTEMS

There are four main functions of ground stations or earth stations in the overall space system. They are tracking, telemetry, control, and data processing and distribution. To ensure that the spacecraft is in its correct orbit, it must be tracked by a ground station. This track information can then be interpreted and navigational corrections provided by an uplink. The active telemetry provided on a downlink from the satellite is diagnosed at the ground station to see that all systems are operating within their bounds. For example, are the thermal control panels providing adequate heat regulation, is there a malfunction in an electronic component, are the solar panels inclined to the correct attitude for maximum absorption of sunlight directed at them? These and other data allow the operator at the ground station to continously monitor the health and welfare of the spacecraft. If any of the information provided on the downlink indicates a difficulty on the platform then corrective measures may be taken; thus, control is provided.

Further control may be exercised by commanding the satellite to dump stored data. Upon receiving the down linked data upon request or at specified intervals, the earth station must process this data into usable form, then distribute it to its consumers or users. In the case of a weather satellite, the data may pertain to various cloud formations and, when processed, the product may be the satellite image of clouds seen on the evening news.

A concluding note on ground systems is that within a single satellite system the number of ground earth stations are dependent upon the altitude of the orbit of the satellite and its mission. Thus, the satellite system is an integrated complicated network from launch to orbit, from mission accomplishment to data processing and to product dissemination. Although, possible sensors aboard a spacecraft were suggested, the missions of spacecraft were not fully described. The next section will discuss a few of the missions of satellite systems.

K. MISSION

Several general mission categories are associated with the satellite payload and each need separate description. These mission categories are; navigation, surveillance, communication, and early warning.

1. NAVIGATION

The requirement for accurate navigation has never been as critical as it is today. Throughout history, the need for accurate positioning has been a concern of sailors, airmen, surveyors, soldiers, etc., but, the requirement for accuracy could be met with the instruments at hand, the stars, sun and horizon. Today, the requirements for accuracy far exceed those of the past in the fields of resource location, global positioning, and particularly in the area of military concerns. In order to present a credible defense, a military unit must be able to know within precise limits (meters) where it is on the surface of the earth. The positioning is the base line for detection, localization, targeting and destruction of hostile forces, and thus positioning is of utmost importance to the U.S. Navy. To this end, the Navy has fielded TRANSIT, a space-based navigational satellite system. This and other systems provide a "stable carrier frequency. . ., continually transmiting a message conveying current satellite time and description of its orbit" [Ref. 52:p. 45]. The atomic clocks used on modern satellites provide extremely accurate times which vary as little as one second every 30,000 year [Ref. 52:p. 44]. Thus, the ability to measure time very precisely, to know the position of the satellites in its orbit, and to detect Doppler shift in the frequencies broadcast by the satellite provides the elements of operation for TRANSIT.

Although TRANSIT met most of the requirements of the Navy since its FOC in 1965, utilizing four satellites in equally spaced polar orbits, new requirements arose which TRANSIT could not meet. Thus, in 1973, a Joint Services program was initiated to provide a system that would be passive, continuous and operational under all conditions, unifying coordinates on land, sea, and air. The Global Position System (GPS) was the outfall of this initiative and it will became fully operational in the 1980's (currently providing limited operations).

GPS uses TDOA from 2-4 satellites. Two satellites when time is accurately known, provides latitude and longitude to the user. The user may obtain latitude, longitude, altitude, and time when four (4) satellites are in his field of view.

2. COMMUNICATIONS

Communications utilizing satellites has always been on the leading edge of satellite technology dating back to the late 1950's. It was obvious, early on, that the ability to communicate world wide via satellite was an attractive one.

The principle of operation is rather straight forward. A signal originating at an earth station is broadcast toward the satellite which the satellite receives and rebroadcasts to another earth station within its field of view. The higher the altitude the greater the satellite's field of view, but the higher the altitude the greater the transmission loss, and hence, the lower the bit rate of transmission, and less the amount of information passed. Transmission loss, may in part, be overcome by antenna design and amplifiers. The culmination of gains and losses is a stream of data with an acceptable error rate, which may be voice, facsimile or video images.

The Navy has used many satellite systems in the past to meet its communication need. These include SYNCOM, LES-6, TACSAT-1, TACINTELL and FLTSATCOM. Fleet Satellite Communication (FLTSATCOM) satellites are an eight foot hexagonal vehicle, fifty inches high, carrying a sixteen foot parabolic dish antenna and weighing 1,860 kilograms [Ref. 52:p. 37]. It provides twenty-three separate channels which provide services to ships, grounds units, and aircraft.

FLTSATCOM will have continual demands placed upon it, there will always be more to communicate than frequency bandwidth upon which to place it.

3. SURVEILLANCE and TARGETING

"Sensors that monitor events on the earth's surface have a direct application of war force. Satellites provide an eye in the sky that can remain there for long periods of time. Thus, they offer considerable potential for monitoring the actions of an enemy in cases where there is no other means to do so. The catalog of possible sensors is long; television or infrared cameras, magnetometers, electronic scanners or receivers, and nuclear detectors all offer possibilities for the enhancement of naval war force." [Ref. 52:p. 5]

> "As information grows older, its tactical value diminishes greatly. Therefore, in order for the data acquired by spacecraft to be useful to the Naval Commander, they must be available to him soon after they are recorded. Thus, the Navy's interest in space includes not only the satellites and their sensors, but the means for transferring the information to the fleet user. The capability for the "real time reactant...the transfer of data as they are acquired...is a primary goal in the development of many space systems." [Ref. 52:p. 5]

> "The potential value of space technology to naval warfare is great. Whether or not the United States exploits this potential to the fullest extent possible, other countries will continue to do so. Prudence demands the best effort we can make". [Ref. 52:p. 5]

4. Early Warning

With the potential satellites for use in the warfare arena as described in the previous section on Surveillance and Targeting, one soon realizes the possibility of their use in an early warning role. Again, the technologies which might prove useful are: television or infrared cameras, magnetometers, electronic scanners or receivers, nuclear detectors, and radar.

VII. The Integration of AI Into Space Systems

It is hoped that at this juncture the reader is comfortable with the fundamental technologies of AI and the fundamental characteristics of space systems. Knowledge of these topics will make the discussion of their integration more obvious. Prior to undertaking such a discussion it is worthwhile to emphasize the justification for the use of AI.

AI offers its greatest benefits by first, proving cost effective in information processing, decision support, and in preserving expertise. Second, it may displace humans from montonous or hazardous situations. Third, it may provide mobility, and thus, survivibility to ground nodes. Mobility may also be an important element of survival to the space vehicle as an autonomous vehicle, if it is threatened. Fourth, AI may provide system reliability by redundancy--redundancy of expertise which may allow for the graceful degradation of a damaged system. These key points should be kept in mind while considering Figure 1 which illustrates in a matrix how AI may be integrated into space systems.

Figure 1 consists of fundamental characteristics of a generic space system, (which is not intended to be an exhaustive list), and the techonolgies of AI. The category of robotics and autonomous vehicles was originally placed in the matrix but was removed when it became readily apparent that it was a collection of technologies represented in columns one, two and three. Thus it is ommitted here.

	Expert Systems	Natural Language Processing	Pattern Recognition
Space Bus Functions			
thermal control	X		
altitude control	X		
attitude control	X		
survivability	X		
power management	X		
Mission Functions			
sensor control	X		
sensor queuing	X		X
target planning	X		
mission planning	X	X	X
data integration	X		X
product production	X	X	X
signal recognition	X		X
Ground Node Functio			
TT & C	X	X	
product production	X	X	X
product distribution	X	X	X
mission planning	X	X	X
signal recognition	X	X	X
Spacecraft Missions			
navigation	X	X	X
surveillance	X	X	X
communication	X	X	X
early warning	X	X	X

Elements of a space system will be addressed with regard to the application of AI; however, where the outcome is obvious or repetitive, that element will be disregarded.

On board processing (OBP) is considered critical and essential to the application of AI to space systems. Its presence gives a great deal of flexibility to the entire system, may reduce both reliance upon ground nodes and time delays within the system. For these reasons OBP will be considered present in the following discussion.

A. SPACE BUS FUNCTIONS

The space bus functions may be viewed as a group of elements which operate to provide control for the well being of the space craft. This type of control function is inherent in thermal, altitude, attitude, and power control systemS. Each of these systems serve as closed-loop monitors. Prescribed thresholds, or out-of-bounds limits, may be set for any given subsystem and appropriate action when the limits are exceeded. The thermal control system is offered as an example:

> As described earlier, the thermal system is responsible for keeping the temperature of the space craft within acceptable limits. These limits may be coded into an ES which would adhere to these boundaries. A simple example of a rule in ES might be:
>
> ```
> if the temperature is greater
>    than X degrees
> and the thermal indicators are
> ```

AD-A174 250 AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND ITS POTENTIAL USE IN SPACE SYSTEMS(U) NAVAL POSTGRADUATE SCHOOL MONTEREY CA G W MCDONALD JUN 86 2/2

UNCLASSIFIED F/G 22/2 NL

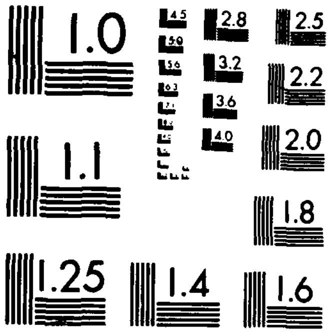

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

```
   operating correctly
then rotate the thermal panels
   Y degrees to the reflective side
```

Granted, this is an oversimplification to what is a complicated or multifarious control problem, but the principles of solution remain the same:

1) set boundaries (rule)

2) monitor the situation (sensors)

3) collect information (data)

4) combine the data with the rule (action)

5) achieve desired state (goal)

6) iterate

Though the rules, sensors, data, subsystems, and specific goals may change, the principle of solution is unchanged; thus, any control function of the space bus, or other control function within the space system may be approached with the same methodology.

B. MISSION FUNCTIONS

Mission functions are data dependent and therefore appropriate for the application of AI. However, the diagnostic and predictive models for these functions are much more complicated then the ones examined on the space bus. As the nature of the problem becomes more complicated the need for a variety of methods for manipulation and presentation of information becomes critical. Natural language processing and pattern recognition become useful tools in managing these complicated ideas.

If sensor control is considered, it will become apparent that scene or pattern recognition is instrumental in problem solving. For example, a weather satellite collects information on cloud formations (pattern) which are characteristic of a hurricane. As the information collected by its sensors supports the presence of the characteristics and these in turn evince a pattern, the pattern may be correlated with the pattern of a hurricane. If single sensor information postulates a pattern but has inconclusive evidence to support a specific conclusion then it may cue another sensor, e.g. a wind speed sensor, to contribute further data to the decision process. Thus, the collaboration of an imaging sensor and a wind speed indicator may provide enough data for the ES to conclude a pattern characteristic of a hurricane. Pattern recognition as applied to a mission function is more fully discussed in the thesis written by Paul Schuh at the Naval Postgraduate School, March 1986, "Applications of Expert Systems Techniques to Classic Wizard Product Patterns (U)."

The area of product production offers an excellent opportunity for natural language processing both on the constructive and interpretive ends. When the product is textual in nature, NRP can be used to construct messages by frame representation for dissemination which can be easily interpreted at the receiving end using the same frame representation. In the case of narrative text, although the problem becomes more difficult, NLP may provide the parser which could be used in interpreting this narrative.

Mission planning also offers an excellent arena in which AI may operate. As mission planning is a resource management problem similar to control, an expert system may prove to be effective in managing its assets. In mission planning the time a sensor is available to support the mission is the asset, and its most efficient use is essential. But how may the demands upon this asset be prioritized and continually regulated? An ES offers the ability to do that by firing rules which would enable sensors to support the spacecraft mission. Pointing of high resolution cameras, such as the space telescope, is a likely candidate for the application of this type of technology. Or, in a weather sensing satellite, sensors may have been directed to search in a specific direction for a specific phenomenon, based upon priorities. However, these priorities may be changed based upon a pattern of clouds which may indicate a more urgent need for the use of the sensors. Thus, priorities for the sensor utilization have been changed by the cooperative efforts of an ES and a program for pattern recognition.

Resource management of the spacecraft is an application for which the ES may be the most beneficial. As another example of this, consider target planning. This is a scheduling function, wh·ch must be optimized. LANDSAT may have a variety of geographic locations to image and an ES can schedule the imaging of these locations, continually updating this scheduling as environmental events dictate change is necessary, e.g., clouds, rain storms, or any event which would significantly degrade LANDSAT's imaging capability.

Many functions at the ground node on data are data intensive and require a human expert to carefully monitor and respond to a change in these functions. The information exchange involved in the TT&C process may be controlled by an ES. Again, it is a matter of monitoring information which indicate whether the spacecraft is operating within prescribed limits. Upon a transgression of a limit, the ES would detect it, and cue a synthesized voice alert. This alert would then make a human monitor aware that a problem exists and allow the human to concentrate his attention on the matter at hand. If the problem is a familiar one the human monitor may elect to allow an ES to offer solutions, or to proceed to enact solutions without over riding the ES.

Although product production and distribution have been described as a possible mission function aboard the spacecraft, it is likely that these functions will, for the time being, remain a function of the ground node, as they require a great deal of computer capability. Thus, the production of products and their routing is an example of the use of AI. Not only the construction and interpretation of messages but the dissemination of messages (products) may be controlled by AI techniques.

Upon considering the missions of navigation, communication, surveillance, and early warning, the attributes of AI technologies and their applications to these missions are easily understood. These missions are a collection of spacecraft bus functions, mission functions, and cooperating ground node entities. Thus, where AI has been applicable to the elements of these subcategories it is also applicable to the system mission function.

Figure 1 represents a functional description of how AI may be applied to space systems. It is a description of what the architecture of this integration might be like. A few general examples are offered but the real work of how this integration is to take place is not discussed. It is this "how to" make the integration happen which is the next step on the ladder to successfully seeing AI and space systems become collaborating systems. The "how to" is the engineering design function which is the doorstep at which this paper ends.

It may seem anticlimactic to end this thesis with a functional description. It has the familiar sound of predictive claims made about AI throughout its history. It is true that the design function of integrating AI and space systems is beyond the scope of this paper and the current capabilities of this author, several conclusions, however, have been reached. The conclusions are presented in the following chapter and are divided into general, and specific prioritized, categories.

VII. CONCLUSION

The following general conclusions are offered regarding AI and its potential application to space systems.

1) AI technologies, with some effort, can be understood and their possible applications recognized.

2) AI systems are difficult to construct.

3) Because of the difficulty of construction, there exists a large gap between the functional description of AI and engineering design of AI.

4) Few people are capable of engineering AI systems but many are needed.

5) Only through the actual engineering of AI systems and their operational applications to space systems, will the assets and liabilities of these combined technologies be realized.

6) The application of AI to space systems will come slowly and should be carefully scrutinized.

7) Currently there is no good, and only in very limited instances is there satisfactory, substitute for human intelligence.

More specifically, AI may be applied to the following functions in a space system. These applications are prioritized given the current state of the art of AI.

A. ENGINEERING CONTROL FUNCTIONS

In this environment, data is generally predictive and expected within a limited domain. Thus, control of temperature, altitude, attitude, power budgeting, etc., are the most likely candidates for the application of AI technologies.

B. DATA ANALYSIS

The correlation and analysis of numerical representation of areas as seemingly diverse as ship hull type to electromagnetic pulse pattern recognition are likely areas for the application of AI. However, data analysis in the case of a large number of variables being more complicated then control functions, require a greater degree of reasoning capability, making them less amenable to current AI technologies.

C. PRODUCT GENERATION

This area represents a variety of challenges. Standard formatted, clearly defined messages may not present overwhelming difficulties for AI. However, as the domain becomes less clearly defined and the format of the product cannot be clearly specified (as in text generation and understanding) the application of AI is currently less appealing, but, not without promise.

D. MISSION PLANNING

Concerning the variables of assets, requirements, and anticipated or predictable changes within a restricted domain, AI technologies may provide the efficient application of assets to requirements.

The obvious next stop is then to take the technologies of AI in whatever stage of maturity they are found and press on. Even though tough problems are presented there is no reason to abandon a new science which continues to present great promise. The case for continued research has seldom been more appropriate than when applied to AI and space systems integration. Regarding this thesis and the matrix in the preceding chapter, the same applies; they represent a starting point from which a more detailed development may take place. The author hopes these conclusions will stimulate further research in the area.

Some of these conclusions are more fully explored in a thesis in preparation by Lt. Debra K. Anderson, Naval Postgraduate School, which is a companion volume to this thesis.

GLOSSARY

This glossary consists of terms from the following works:

(1) Tennant, H., Natural Language Processing, PBI, New York, NY, 1981

(2) Waterman, D., A Guide to Expert Systems, Addison-Wesley Publishing Company, Reading, MA, 1986.

(3) Hayes-Roth, F., Waterman, D., and Lonat, D., Building Expert Systems, Vol. 1, Addison-Wesley Publishing Company, Reading, MA, 1983.

In some instances two definitions are offered for the same term, in order to more clearly develop the idea presented by the term.

access-oriented methods
Programming methods based on the use of probes that trigger new computations when data are changed or read. (2)

active value
A procedure invoked when program data are changed or read, often used to drive graphical displays of gauges that show the values of the program variables. (2)

agenda
A prioritized list of pending activities, usually the applications of various pieces of knowledge. (2)

AI
Artificial intelligence. (2)

algorithm
A formal procedure guaranteed to produce correct or optimal solutions. (2)

Artificial intelligence
The subfield of computer science concerned with developing intelligent computer programs. This includes programs that can solve problems, learn from experience, understand language, interpret visual scenes, and, in general, behave in a way that would be considered intelligent if observed in a human (2)

ATN-augmented transition network
A heavily used parsing formalism composed of a grammar that is applied in a recursive, top down fashion, and augmented with global registers that are capable of temporarily holding structures for latter use. (1)

back-chaining
A control procedure that attempts to achieve goals recursively, first by enumerating antecedents that would be sufficient for goal attainment and, second, by attempting to achieve or establish the antecedents themselves as goals. (also backward-chaining) (3)

backward-chaining
An inference method where the system starts with what it wants to prove, e.g., Z-, and tries to establish the facts it needs to prove Z. The facts needed to prove a conjecture (Z) are typically given in rule form; e.g., IF A & B, THEN Z. If A and B aren't known (aren't available as data), the system will try to prove A and B by establishing any additional facts (as specified by other rules) needed to prove them. The additional facts are established the same way A and B were established, and the process continues until all needed facts are established or the system gives up in defeat. (2)

backtracking
A search procedure that makes guesses at various points during problem-solving, returning to a previous point to make another choice when a guess leads to an unacceptable result. (3)

belief
(1) A hypothesis about some unobservable situation. (2) A measure of the believer's confidence in an uncertain proposition. (3)

blackboard
A data base accessible to independent knowledge sources and used by them to communicate with one another. The information they provide each other consists primarily of intermediate results of problem solving. (2)

blackboard architecture
A way of representing and controlling knowledge based on using independent groups of rules called knowledge sources that communicate through a central data base called a blackboard. (2)

bottom-up parsing
Synonymous with data-driven parsing; a parsing method that starts with the lowest structure (e.g., words) and builds higher level structures from them (e.g., noun phrases, prepositional phrases, sentences) (1)

break package
A mechanism in a programming or knowledge engineering language for telling the program where to stop so the programmer can examine the values of variables at that point. (2)

C
A low-level, efficient, general-purpose programming language associated with the UNIX operating system. C is normally used for system programming. (2)

CAD Computer-aided design
The use of computer technology to assist in the design process, e.g., the design of integrated circuits. (2)

CAI
Computer-assisted instruction; the application of computers to education. The computer monitors and controls the student's learning, adjusting its presentation based on the responses of the student. (2)

certainty factor
A number that measures the certainty or confidence one has that a fact or rule is valid. (2)

conflict resolution
The technique of resolving the problem of multiple matches in a rule-based system. When more than one rule's antecedent matches the data base, a conflict arises since (1) every matched rule could appropriately be executed next, and (2) only one rule can actually be executed next. A common conflict resolution method is priority ordering, where each rule has an assigned priority and the highest priority rule that currently matches the data base is executed next. (2)

cooperating knowledge sources
Specialized modules in an expert system that independently analyze the data and communicate via a central, structured data base called a blackboard. (2)

data base
The set of facts, assertions, and conclusions used to match against the IF-part of rules in a rule-based system. (2)

dependency
A relation between the antecedents and corresponding consequents produced as a result of applying an inferential rule. Dependencies provide a record of the manner in which decisions are derived from prior data and decisions. (3)

dependency-directed backtracking
A programming technique that allows a system to remove the effects of incorrect assumptions during its search for a solution to a problem. As the system infers new information, it keep dependency records of all its deductions and assumptions, showing how they were derived. When the system finds that an assumption was incorrect, it backtracks through the chains of inferences, removing conclusions based on the faulty assumption. (2)

domain expert
A person who, through years of training and experience, has become extremely proficient at problem solving in a particular domain. (2)

domain knowledge
Knowledge about the problem domain, e.g., knowledge about geology in an expert system for finding mineral deposits. (2)

ellipsis
The omission of words or phrases in an utterance, with the assumption that the listener can use the current context to assume what has been omitted. (1)

evaluation function
A procedure used to determine the value or worth of proposed intermediate steps during a hunt through a search space for a solution to a problem. (2)

evolutionary development (of software)
The practice of iteratively designing, implementing, evaluating, and refining computer applications, especially characteristic of the process of building expert systems. (3)

exhaustive search
A problem-solving technique in which the problem solver systematically tries all possible solutions in some "brute force" manner until it finds an acceptable one. (2)

expectation-driven reasoning
A control procedure that employs current data and decisions to formulate hypotheses about yet unobserved events and to allocate resources to activities that confirm, disconfirm, or monitor the expected events. (3)

exophoric
Depending on the external (non-linguistic) situation for interpretation, as in "Did you hear that explosion?" following an explosion. (1)

expert system
A computer program that uses expert knowledge to attain high levels of performance in a narrow problem area. These programs typically represent knowledge symbolically, examine and explain their reasoning processes, and address problem areas that require years of special training and education for humans to master. (2)

expert-system-building tool
The programming language and support package used to build the expert system. (2)

expertise
The set of capabilities that underlies the high performance of human experts, including extensive domain knowledge, heuristic rules that simplify and improve approaches to problem-solving, metaknowledge and metacognition, and compiled forms of behavior that afford great economy in skilled performance. (3)

explanation facility
That part of an expert system that explains how solutions were reached and justifies the steps used to reach them. (2)

facet
See slot. (3)

fact
A proposition or datum whose validity is accepted. (3)

frame

A knowledge cluster that embodies what an individual knows about one particular concept; a frame system is an individual's knowledge about the world represented by frame. (1)

frame

A knowledge representation scheme that associates one or more features with an object in terms of various slots and particular slot-values. Similar to property-list, schema, unit, and record in various writings. (3)

frame

A knowledge representation method that associates features with nodes representing concepts or objects. The features are described in terms of attributes (called slots) and their values. The nodes form a network connected by relations and organized into a hierarcy. Each node's slots can be filled with values to help describe the concept that the node represents. The process of adding or removing values from the slots can activate procedures (self-contained pieces of code) attached to the slots. These procedures may then modify values in other slots, continuing the process until the desired goal is achieved. (2)

frame-based methods

Programming methods using frame hierarchies for inheritance and procedural attachment. (2)

forward chaining

An inference method where the IF-portion of rules are matched against facts to establish new facts. (2)

forward chaining

A control procedure that produces new decisions recursively by affirming the consequent propositions associated within an inferential rule with antecedent conditions that are currently believed. As new affirmed propositions change the current set of beliefs, additional rules are applied recursively. (3)

fuzzy logic

An approach to approximate reasoning in which truth values and quantifiers are defined as possibility distributions that carry linguistic labels, such as true, very true, not very true, many, not very many, few and several. The rules of inference are approximate, rather than exact, in order to better manipulate information that is incomplete, imprecise, or unreliable. (2)

garden path sentences
Sentences that generally force listeners to consciously back up and reinterpret them, such as "I was wary of Ali's punch, but by the third round I realized there was no liquor in it"; sentences that people apparently parse non-deterministically. (1)

generality hierarchy
A tree structure of concepts, where the most general concepts are closest to the root, and most specific closest to the leaves; more specific concepts generally inherit the characteristics of their ancestors. (1)

general-purpose knowledge engineering language
A computer language designed for building expert systems and incorporating features that make it applicable to different problem areas and types. (2)

generate and test
A problem-solving technique involving a generator that produces possible solutions and an evaluator that tests the acceptability of those solutions. (2)

goal-directed reasoning
See back-chaining. (3)

HEARSAY-II architecture
The organization of a problem-solving system in terms of several cooperating, independent specialists representing diverse areas of knowledge, which exchange partial results via a blackboard and collectively assemble an overall solution incrementally and opportunistically. (3)

heuristic
A rule of thumb or simplification that limits the search for solutions in domains that are difficult and poorly understood. (1)

heuristic parsing
Technique (generally associated with ATN parsers) of ordering the hypotheses in a top-down parser (acts in an ATN parser) to try the most likely first, in the hope that the first parse found is the most likely to be correct. (1)

heuristic programming project
The research group at Stanford University that principally pioneered the field of knowledge engineering and produced the largest collection of expert systems. (2)

heuristic rule
A procedural tip or incomplete method for performing some task. (3)

human engineering
(A misnomer). The task of designing human-machine interfaces to achieve effective human utilization of machine capacities. (3)

hypothetical worlds
A way of structuring knowledge in a knowledge-based system that defines the contexts (hypothetical worlds) in which facts and rules apply. (2)

ICAI
Intelligent computer-assisted instruction; the application of AI methods to the CAI problem. (2)

image understanding
The use of AI methods to process and interpret visual images, e.g., analyzing the signals produced by a TV camera to recognize and classify the types of objects in the picture. (2)

inference chain
The sequence of steps or rule applications used by a rule-based system to reach a conclusion. (2)

inference, data-directed
See **forward-chaining**

inference engine
That part of a knowledge-bases sytem or expert system that contains the general problem-solving knowledge. The inference engine processes the domain knowledge (located in the knowledge base) to reach new conclusions. (2)

inference net
All possible inference chains that can be generated from the rules in a rule-based system. (2)

inferential rule
An association between antecedent conditions and consequent beliefs that enables the consequent beliefs to be inferred (deduced) from valid antecedent conditions. (3)

inheritance hierarchy
A structure in semantic net or frame system that permits items lower in the net to inherit properities from items

instantiation
An object that fits the general description of some class or, specifically, a pending process that associates specific data objects with the parameters of a general procedure. (3)

INTERLISP
An elaborate programming system providing extensive programming support for constructing and maintaing large LISP programs. (3)

interpreter
In an expert system, that part of the inference engine that decides how to apply the domain knowledge. In a programming system, tnat part of the system that analyzes the code to decide what actions to take next. (2)

I/O
Input/output; the communication between a computer program and its user. (2)

KE
Knowledge engineer. (2)

knowledge
The information a computer program must have to behave intelligently. (2)

knowledge acquisition
The process of extracting, structuring, and organizing knowledge from some source, usually human experts, so it can be used in a program. (2)

knowledge base
The portion of a knowledge-based system or expert system that containes the domain knowledge. (2)

knowledge-based system
A program in which the domain knowledge is explicit and separate from the program's other knowledge. (2)

knowledge engineering
The process of building expert systems. (2)

knowledge representation
The process of structuring knowledge about a problem in a way tnat make the problem easier to solve. (2)

LCD
Liquid crystal display. (2)

LISP
The principal programming language of AI, which provides an elegant, recursive, untyped, and applicative framework for symbolic computing; actually a family of variants. (3)

list structure
A collection of items enclosed by parentheses, where each item can be either a symbol or another list, e.g., (ENGINE FUEL (Y5 BILLO 23 (CLAY 7)). (2)

logic-based methods
Programming methods that use predicate calculus to structure the program and guide execution. (2)

LSI
Large scale integration. See VLSI. (2)

MACLISP
The variant of LISP developed and promulgated by workers at MIT. (3)

meta
Prefix designating reflexive applications of the associated concept. (3)

metacognition
The capability to think about one's own thought processes. (3)

metaknowledge
Knowledge about knowledge. (3)

metaknowledge
Knowledge in an expert system about how the system operates or reasons, such as knowledge about the use and control of domain knowledge. More generally, knowledge about knowledge. (2)

metalevel knowledge
See **metaknowledge.** (2)

metarule
A rule that describes how other rules should be used or modified. (2)

multiple lines of reasoning
A problem-solving technique in which a limited number of possibly independent approaches to solving the problem are developed in parallel. (2)

natural language
The conventional method for exchanging information between people, such as English as a means of communication for human speakers and various formal written systems as a means of representing intentions in technical disciplines (chemical graphs, DNA sequences, engineering diagrams, and so on). (3)

non-deterministic parsing
parsing which allows decisions to be changed or allows several alternative interpretations to proceed in parallel (see deterministic parsing). (1)

nonmonotonic reasoning
A reasoning technique that supports multiple lines of reasoning (multiple ways to reach the same conclusion) and the retraction of facts or conclusions, given new information. It is useful for processing unreliable knowledge and data. (2)

object-oriented methods
Programming methods based on the use of items called **objects** that communicate with one another via messages in the form of global broadcasts. (2)

parser
Generally intended as a formalism that assigns a structural description to a sentence; also used to describe formalisms that assign a semantic interpretation to a sentence (as in a parser for a semantic grammar). (1)

perlocutionary acts
The effects that a speaker actually has on a listener. (1)

phonetic structure
Shows the structure of a sentence as it would actually be pronounced (see surface structure, deep structure, conceptual structure). (1)

phrase structure grammar
A set of rules that indicate what and how categories of words and phrases can be combined to construct other categories of phrases. (1)

pragmatic ellipsis
An omission of information from a syntactically complete sentence that must be assumed from the context. (1)

pragmatics
The study of the role of contextual knowledge in language; knowledge about the world. (1)

predicate calculus
A formal language of classical logic that uses functions and predicates to describe relations between individual entities. (2)

probability propagation
The adjusting of probabilities at the nodes in an inference net to account for the effect of new information about the probability at a particular node. (2)

problem-oriented language
A computer language designed for a particular class of problems, e.g., FORTRAN designed for efficiently performing algebraic computations and COBOL with features for business record keeping. (2)

problem reformulation
Converting a problem stated in some arbitrary way to a form that lends itself to a fast, efficient solutions. (2)

problem space
See **search space.** (2)

problem-solving methods, weak and strong
Heuristic for control. Weak methods are domain independent, while strong methods exploit domain knowledge to achieve greater performance. (3)

procedure-oriented methods
Programming methods using nested subroutines to organize and control program execution. (2)

production
An IT-THEN statement or rule used to represent knowledge in a human's long-term memory. (2)

production rule
The type of rule used in a production system, usually expressed as IF **condition** THEN **action.** (2)

production system
A type of rule-based system containing IF-THEN statements with conditions that may be satisfied in a data base and actions that may change the data base. (2)

property
See **slot.** (3)

property list
A construct in LISP that associates with an object called an atom a set of one or more pairs, each composed of a "property" and a "value" of that property for that object. (3)

pruning
Reducing or narrowing the alternatives, normally used in the context of reducing possibilities in a branching **tree structure** such as the search through a problem space. (2)

real-world problem
A complex, practical problem which has a solution that is useful in some cost-effective way. (2)

representation
The process of formulating or viewing a problem so it will be easy to solve. (2)

resolution theorem proving
A particular use of deductive logic for proving theorems in the first-order predicate calculus. The method makes use of the following resolution principle: (A v B) and (-A v C) implies (B v C). (2)

robustness
That quality of a problem solver that permits a gradual degradation in performance when it is pushed to the limits of its scope of expertise if given errorful, inconsistent, or incomplete data or rules. (2)

rule
A formal way of specifying a recommendation, directive, or strategy, expressed as IF **promise** THEN **conclusion** or IF **condition** THEN **action.** (2)

rule
A pair, composed of an antecedent conditon and a consequent propositon, which can support deductive processes such as back-chaining and forward-chaining. (See also **heuristic rule.**) (2)

rule-based methods
Programming methods using IF-THEN rules to perform forward or backward chaining. (2)

rule-based program

A computer program that explicitly incorporates rules or ruleset components. (3)

ruleset

A collection of rules that constitutes a module of heuristic knowledge. (3)

satisfice

Achieve a solution that satisfies all imposing constraints. (Opposed to "optimize".) (3)

scaling problem

The difficulty associated with trying to apply problem-solving techniques developed for a simplified version of a problem to the actual problem itself. (2)

scene analysis

See **image analysis.** (2)

scheduler

The part of the inference engine that decides when and in what order to apply different pieces of domain knowledge. (2)

schema

Conceptual structures equivalent to frames. (1)

scheduling

Determining the order of activities for execution, usually based on control heuristics. (See also **agenda.**) (3)

scripts

Conceptual structures that describe events and sequences of events. (1)

search

The process of looking through the set of possible solutions to a problem in order to find an acceptable solution. (2)

search space

The set of all possible solutions to a problem. (2)

semantic

Pertaining to the meaning, intention, or significance of a symbolic expression, as opposed to its form. (Contrast **syntactic.**) (3)

semantic grammar
A grammar which parses according to semantic categories of words and phrases rather than syntactic categories. (1)

semantic marker
An attribute assigned to a work or phrase indicating that it describes a concept of a particular semantic class; used as restrictions for the selection of competing semantic interpretations. (1)

semantic net
A network representation of knowledge; a broad range of knowledge representation formalisms have been called semantic nets, they all involve the designation of conceptual entities linked to one another by name relations. (1)

semantic net
A knowledge representation method consisting of a network of nodes, standing for concepts or objects, connected by arcs describing the relations between the nodes. (2)

semantics
The study of the relationship between symbols and their meaning; sometimes called what is left of linguistics without syntax, most consider it to be the link between syntax and knowledge representation. (1)

skeletal knowledge engineering language
A computer language designed for building expert systems and derived by removing all domain-specific knowledge from an existing expert system. (2)

skeletal system
See **skeletal knowledge engineering language.** (2)

skill
The efficient and effective application of knowledge to produce solutions in some problem domain. (2)

slot
An attribute associated with a node in a frame system. The node may stand for an object, concept, or event; e.g., a node representing the object **employee** might have a slot for the attribute **name** and one for the attribute **address.** These slots would then be filled with the employee's actual name and address. (2)

slot
A feature or component description of an object in a frame. Slots may correspond to intrinsic features such as name, definition, or creator; or may represent derived attributes such as value, significance, or analogous objects. (3)

specialist
An expert in a narrow problem domain, especially one of the several expert subsystems that cooperate in a HEARSAY-II architecture. (3)

speech understanding
The use of AI methods to process and interpret audio signals representing human speech. (2)

stonewalling
Giving literal but not complete answers to questions; a characteristic of many question answering systems. (1)

story grammar
A grammar describing the allowable structure of stories. (1)

substitution
Anaphoric replacement of a work or phrase by a substitute word or phrase like "the green one" for "the green volleyball"; closely relates to ellipsis. (1)

support environment
Facilities associated with an expert-system-building tool that help the user interact with the expert system. These may include sophisticated debugging aids, friendly editing programs, and advanced graphic devices. (2)

support facilities
See **support environment.** (2)

symbol
A string of characters that stands for some real-world concept. (2)

symbol-manipulation language
A computer language designed expressly for representing and manipulating complex concepts, e.g., LISP and PROLOG. (2)

symbolic reasoning
Problem solving based on the application of strategies and heuristics to manipulate symbols standing for problem concepts. (2)

syntactic
Pertaining to the form or structure of a symbolic expression, as opposed to its meaining or significance. (Contrast **semantic.**) (3)

syntax
The structural description of a language. (1)

tool
A shorthand notation for expert-system-building tool. (2)

tool builder
The person who designs and builds the expert-system-building tool. (2)

tools for knowledge engineering
Programming systems that simplify expert system development. They include languages, programs, and facilities that assist the knowledge engineer. (2)

tools for knowledge engineering
Programming systems that simplify the work of building expert systems, especially generic task packages such as EMYCIN and very high-level languages for heuristic programming such as ROSIE. (3)

top-down parsing
Synonymous with hypothesis driven parsing, similar to expectation driven parsing; a parsing method that hypothesizes a high level structure (e.g. a sentence), then attempts to match (recursively) lower level structures to it. (1)

toy problem
An artificial problem, such as a game, or an unrealistic adaptation of a complex problem. (2)

tracing facility
A mechanism in a programming or knowledge engineering language that can display the rules or subroutines executed, including the values of variables used. (2)

transformational grammar
A theory of syntax that describes the structure of sentences in the language with a set of recursive rules (transformations) that relate pairs of tree structures to one another. (1)

tree structure
A way of organizing information as a connected graph where each node can branch into other nodes deeper in the structure. (2)

truth maintenance
(A misnomer). The task of preserving consistent beliefs in a reasoning system whose beliefs change over time. (3)

underlying representation
The internalized representation or meaning of an utterance; like deep structure and conceptual structure, but does not imply commitment to a particular theory of semantics. (1)

units
A frame-like representation formalism employing slots with values and procedures attached to them. (2)

user
A person who uses an expert system, such as an end-user, a domain expert, a knowledge engineer, a tool builder, or a clerical staff member. (2)

VLSI
Very large scale integration; the development of complex and powerful circuits on small chips. (2)

well-formed substrings
Substrings that are found through the course of parsing that need not be reparsed in the event of backtracking. (1)

REFERENCES

1. Manual, T., and Evanczuk, S., "Artificial Intelligence: Commerical Products Begin to Emerge From Decades of Research", Electronics, Vol. 56, No. 22, pp.127-31, 3 Nov 83.

2. Lattimore, R., Homer's ILIAD, Machines Who Think, W. H. Freeman and Co., San Francisco, CA, 1979.

3. Asimov, I., "I, Robot", Machines Who Think, W. H. Freeman and Co., San Francisco, CA, 1979.

4. McCorduck, P., Machines Who Think, W. H. Freeman and Co., San Francisco, CA, 1979.

5. Turing, H. M., "Computing Machinery and Intelligence", Minds and Machines, Prentice-Hall Inc., Englewood Cliffs, NJ, 1964.

6. Dreyfus, H., What Computers Can't Do, Harper and Row, NY, 1972.

7. Anderson, R., Minds and Machines, Prentice-Hall Inc., Englewood Cliffs, NJ, 1964.

8. Ernst, G., and Newell, A., GPS: A Case Study In Generality and Problem Solving, Academic Press, NY, 1968.

9. Charniak, E., and McDermitt, D., Artificial Intelligence, Addison-Wesley, Reading, MA, 1985.

10. Dreyfus, H., and Dreyfus, S., "Why Computers May Never Think Like People", Technology Review, Jan 1986.

11. "Artificial Intelligence: An Overview", Defence Science and Techonology, Nov 1985.

12. Feigenbaum, C., and Feldman, J., Computers and Thought, Mcgraw-Hill Co., Inc., NY, 1963.

13. Scriven, M., "The Mechanical Concept of the Mind", Minds and Machines, Prentice-Hall, Englewood Cliffs, NJ, 1964.

14. Foget, L., Owens, A., and Feldman, M., Artificial Intelligence Through Simulated Evolution, J. Wiley & Sons, Inc., NY, 1966.

15. Winston, P., "Perpective", A.I. Business, MIT Press, Cambridge. MA, 1984.

16. McKibbin, W., "Awaiting the Intelligent Computer", INFOSYSTEMS, Vol. 30, No. 8, pp. 96-98, Aug 83.

17. Waltz, D., "Artificial Intelligence", Scientific American, Vol. 247, No. 4, pp. 118-131, Oct 82.

18. Rich, E., Artificial Intelligence, Mcgraw-Hill, NY, 1983.

19. Gevarter, W., Artifical Intelligence, Expert Systems, Computer Vision, and Natural Language Processing, Noyes Pub., Park Ridge, NJ, 1984.

20. Waterman, D., A Guide To Expert Systems, Addison-Wesley Pub. Co., Reading MA, 1986.

21. Davis, R., "Amplifying Expertise with Expert Systems", A.I. Business, MIT Press, Cambridge, MA, 1984.

22. Minsky, M., "The Problems and Promise", A.I. Business, MIT Press, Cambridge, MA, 1984.

23. Schank, R., "Intelligent Advisory Systems", A.I. Business, MIT Press, Cambridge, MA, 1984

24. Baker, J., "Dipmeter Advisor: An Expert Log Analysis System Schlumberger", A.I. Business, MIT Press, Cambridge, MA, 1984.

25. McCune, B., and Drazovich, R., "Radar with Sight and Knowledge", Defense Electronics, Vol. 15, No. 8, pp. 71-79, Aug 1983

26. Meyrowitz, A., "Military Applications of Artifical Intelligence", Signal, Vol. 30, No., 40, pp. 45-49, Jun 1984.

27. Reitman, W., "Artificial Intelligence Applications for Business", Computer and People, Vol. 33, No 7&8, pp. 8-11, Jul-Aug 1983.

28. Hayes-Roth, E., "Knowledge-based Expert Systems - The State of the Art in the U.S.", Expert Systems, A. Wheaton and Co., Limited, Exeter, Devonshire, England, 1984.

29. Hayes-Roth, E., Waterman, E., and Lenat, D., Building Expert Systems, Vol. 1, Addison-Wesley Publishing Co., Inc., Reading, MA, 1983.

30. Brachman, R., and others, "What are Expert Systems", Building Expert Systems, Vol. 1, Addison-Wesley Publishing Col., Inc., Reading MA, 1983.

31. Stefik, M., and others, "The Archeticture of Expert Systems", Building Expert Systems, Addison-Wesley Publishing Co., Inc., Reading, MA 1983.

32. Feigenbaum, A., "Knowledge Engineering: The Applied Side", Intelligent Systems, J. Wiley and Sons, NY, 1983.

33. Weiss, S., and Kalikowski, C., A Practical Guide to Designing Expert Systems, Rowman and Allenheld, Totowa, NJ, 1984.

34. Quinlan, J., "Fundamentals of the Knowledge Engineering Problem", Introductory Readings in Expert Systems, Gordan and Breach Science Pub., NY, 1982.

35. Gevarter, W., Intelligent Machines, Prentice-Hall, Englewood Cliffs, NJ, 1985.

36. Tennant, H., Natural Language Processing, PBI, NY, 1981.

37. Webster's New World Dictionary, World Publishing Co., Cleveland, OH, 1968.

38. Webster, B., "Discourse Model Synthesis Preliminaries to Reference", Natural Language Processing, PBI, NY, 1981.

39. Watanabe, S., Pattern Recognition, J. Wiley and sons, NY, 1985.

40. Marr, D., "Representing and Computing Visual Information", Artificial Intelligence: An MIT Perspective, Vol. 2, MIT Press, Cambridge, MA, 1979.

41. Duff, M., "Seeing Machines", Intelligent Systems, J. Wiley and Sons, NY, 1983.

42. Freitas, R., "The Birth of the Cyborg", ROBOTICS, Anchor Press/Doubleday, Garden City, NY, 1985.

43. Binford, T., "The Machine Sees", ROBOTICS, Anchor Press/Dougleday, Garden City, NY, 1985.

44. Moravec, H., "The Rovers", ROBOTICS, Anchor Press/-Doubleday, Garden City, NY, 1985.

45. Orlando, N., An Intelligent Robotics Control Scheme, paper presented of American Controls Conf., San Diego, CA, 6-8 June 1984.

46. Harrison, W., and Orlando, N., A System-Level Approach to Automation Research, NASA, Langley, VA.

47. Orlando, N., A Space for Intelligent Teleoperated Research; paper presented, AIAA, Harford, CT 25 Oct 1983.

48. Harmon, S., Atonomous Vehicles, Naval Ocean System Center, San Diego, CA.

49. Vogel, J., The Pocket Rocket Reader, Chemical Sys. Div., United Tech.

50. Hazard, H., US Space Launch Systems, Report No. NSSA-R-20-72-2, Navy Space Systems Command, Los Angeles, CA, 1973.

51. Fuhs, A., Spacecraft Design, Notes, Naval Postgraduate School, Monterey, CA, 21 Nov 1985.

52. CNO, Navy Space Digest, CNO, Dept. of Navy, Washington, D.C.

INITIAL DISTRIBUTION LIST

		No. Copies
1.	Defense Technical Information Center Cameron Station Alexandria, Virginia 22304-6145	2
2.	Library, Code 0142 Naval Postgraduate School Monterey, California 93943-5000	2
3.	Herschel H. Loomis, Code 62LM Department of Electrical & Computer Engineering Naval Postgraduate School Monterey, California 93943-5000	1
4.	Elliot H. Weinberg, Code 61WI Physics Department Naval Postgraduate School Monterey, California 93943-5000	1
5.	Lois Brunner, Code 74 C^3 Academic Group Naval Postgraduate School Monterey, California 93943-5000	1
6.	Gary McDonald, Lt., USN HQ US Space Command - J2 Cheyene Mtn. Complex, Colorado 80914	5
7.	Vice Commander United States Space Command Peterson AFB, Colorado 80914-5001	1
8.	FL 2058 Space Library HQ AF SPACECOM/MPSL Peterson Air Force Base, Colorado 80914-5001	1
9.	Mark T. Sandvigen 7538 Wagner Creek Road Talent, Oregon 97540	1
10.	Commander Naval Space Command Dahlgren, VA 22448-5170	1

11. Allen E. Fuhs, Code 72 1
Chairman, Space Systems Academic Group
Naval Postgraduate School
Monterey, California 93943-5000

END

12-86

DTIC